non-specialists in rational and social choice alike. I will want to keep a very keen eye on this handbook – and so do you.

Conrad Heilmann
Erasmus University Rotterdam

doi:10.1017/S0266267112000065

Persons, Interests, and Justice, Nils Holtug, Oxford University Press, 2010, 356 pages.

Scandinavian moral philosophers have become widely recognized as leaders in recent research regarding axiology in general and in broadly defined consequentialism in particular. There are several distinctive features that define their research. Their works tend to be precise, careful, no-nonsense and thought provoking. Additionally, many Scandinavian moral philosophers use the results and formal method of economic theory to elucidate the structure and scope of philosophical arguments. It is no surprise that some of the most exciting results in recent moral philosophy come from this region due to this fruitful combination of economics and philosophy. Nils Holtug is one of the leading moral philosophers in this tradition, and his long-awaited monograph, *Persons, Interests, and Justice*, is a prime example of Scandinavian axiology at its best. Holtug has been working on the fundamental problems posed by Derek Parfit in the last three decades, including personal identity, rationality, population ethics, distributive justice and so on. He has already proposed several innovative arguments about these topics. In *Persons, Interests, and Justice*, Holtug attempts to provide a coherent structure to his previously published works, and put forward a comprehensive theory of justice that meets Parfit's problems. This is a highly ambitious project, and we ought to applaud his courage and motivation.

The book has two parts. The first part is devoted to elucidating 'what matters' in moral theory. As I understand it, by 'what matters', Holtug means what should be understood as the most fundamental building block for morality. According to him, it is self-interest. Roughly speaking, to say that X is in a person's self-interest is to say that X will make her life better and increase her welfare. Throughout the first part, Holtug attempts to identify the most plausible notion of self-interest and its theoretical scope. His favoured notion of self-interest is what he calls the *Prudential View*. The Prudential View is such that a person has a present self-interest in a future benefit if and only if she stands in a relation of continuous physical realization of psychology to the beneficiary. On this view, self-interest is future-looking. This means that the notion of self-interest does not include anything backward-looking, such as concern for the reputation of my deceased grandfather. Once we start talking about the

notion of future benefit, we are led to discussing one of the most difficult questions in moral philosophy: the issue of personal identity. Holtug argues that what matters in the notion of self-interest is not personal identity per se. Rather, he argues that what really matters is the continuous physical realization of psychology, which he calls the *Relation M*. Many sentient non-human animals have the continuous physical realization of psychology, and hence, their welfare should be taken seriously. On the other hand, a foetus does not have significant consciousness (although it is physically continuous to a conscious person) and therefore, does not have relevant self-interest.

The second part is devoted to distributive justice and population ethics. With the exception of Chapter 10, the second part can be read independently of the first. Holtug confines his discussion to persons, and does not include non-human animals for the sake of simplicity. He starts with one important condition that he believes is required for any plausible distributive principle: the *Person-Affecting View*. More specifically, he wants any plausible distributive principle to satisfy *Person-Affecting Outcome Welfarism*, according to which the value of an outcome is an increasing function only of the individual benefits it contains. This condition is welfarist because it has distributive judgements depend only on people's welfare. It is a Person-Affecting View because the overall goodness of outcomes is a strictly increasing function of people's welfare only. Given Person-Affecting Outcome Welfarism, his preferred distributive principle is prioritarianism. He carefully examines (telic) egalitarianism and sufficientarianism, and puts forward many innovative arguments against these. He rejects egalitarianism because it is susceptible to the Levelling Down Objection. He argues that prioritarianism is more acceptable than sufficientarianism because sufficientarianism assigns no positive value to the increase in welfare above a threshold level, which has many counterintuitive implications. What about utilitarianism? Holtug swiftly rejects it on the basis that it is distribution-insensitive. Although he supports prioritarianism in the cases of fixed population size, he admits that prioritarianism, just like any other distributive principle, implies the Repugnant Conclusion in cases where the population size varies. Holtug is very candid about the limit of his preferred prioritarianism, and I like his honesty.

I am impressed and persuaded by Holtug's highly sophisticated arguments in this book. I have no substantive disagreements. However, for the sake of discussion, I will focus on two points that I did not find entirely persuasive. The first point concerns the value of coming into existence in Chapter 5. On the basis of his account of self-interest, Holtug argues that coming into existence can benefit or harm a person. More precisely, a person is benefited by coming into existence if, on balance, his life is worth living, and harmed if, on balance, his life is not worth living.

This applies not only to actual people, but also to merely possible people, that is, people who could have existed but will in fact not. His argument is based on a comparative judgement between existence and non-existence. Holtug contends that existence can be intrinsically better (or worse) for a person than never existing. He appeals to the notion of self-regarding and rational preference in order to ground the comparative judgement of existence and non-existence. According to Holtug, if Jeremy prefers existing to never existing, he has benefited from coming into existence and hence his existence is better for him than non-existence. If he would have preferred to have never existed, then his existence is harmful to him and worse for him than having never existed. I would hope that nobody actually has the latter type of preference. But for present purposes, let us put my hope aside. I have a trouble with his argument. I do not have a clear intuition about the normative comparison about my non-existence. I can proudly report that I actually prefer existing to never existing. But I am not sure whether it follows that my existence is *better* for me than never existing. I do not have a clear intuition concerning whether I am taller than non-existence, ø. Likewise, I do not have a clear intuition concerning whether my existence is better for me than ø. There are two separate issues here. The first issue is whether self-regarding and rational preference best tracks our normative judgement. This is highly debatable. The second and more important issue is whether it makes sense to judge the relative goodness of existence and non-existence. Holtug claims:

> Let us suppose that Jeremy's life contains a net surplus of positive values ... If Jeremy never exists, no positive or negative values accrue to him, and so his non-existence has no value for him. On the basis of these value assessments, existence seems to be better. After all, it seems to be more desirable for a person to have a surplus of positive value than to have no value accrue to him ... (p. 134)

I agree that Jeremy's non-existence has no value for him. However, I do not think that no value is equivalent to zero value. We must distinguish 0 and ø. I take the value of non-existence to be ø. On the other hand, Holtug seems to take it to be zero, and this is why he thinks positive value in Jeremy's existence is greater than no value in his non-existence. Zero value may well be the level worth living. But I do not think that the 'no value' of non-existence is best represented by zero.

The second main point I want to address concerns his criticism of egalitarianism. Holtug takes Parfit's Levelling Down Objection seriously. The Objection runs as follows. Suppose that we lower the level of a better-off person to the level of a worse-off person without benefiting any person. If equality is intrinsically valuable, this levelling-down is, at least in one respect, strictly better. But it is not better in any respect. Therefore, egalitarianism appears to be absurd. This is the Levelling Down Objection

to egalitarianism. It is true that egalitarians hold the Principle of Equality, according to which it is in itself bad if some people are better off than others. It is also true that almost all egalitarians hold the Principle of Utility, according to which it is in itself better if people are better off. As Holtug rightly points out, the Levelling Down Objection does not claim that egalitarians must judge the levelling-down to be, all things considered, better, but that egalitarians must judge the levelling-down to be, at least in one respect, better. If egalitarians hold both the Principle of Equality and Principle of Utility and if they claim that the levelling-down is, all things considered, strictly worse, what is the big deal about claiming that the levelling-down is better in one respect? For instance, consider a version of egalitarianism, where the total goodness of the outcome in the two-person case, G, is given by $G = 1/2(w_1+w_2) - 1/4|w_1-w_2|$, where w_1 and w_2 represent the levels of welfare of the two individuals. According to this, the goodness of the outcome is given by the average welfare and the disvalue of inequality. This formula is equivalent to $1/4 \cdot w_1+3/4 \cdot w_2$ if $w_1>w_2$ and $3/4 \cdot w_1+1/4 \cdot w_2$ otherwise. According to this extensionally equivalent representation, the levelling-down is not better in any respect. In order for the Principle of Equality to be consistent with the Principle of Utility, a suitably small weight ought to be given to the absolute difference between the well-being of two persons. But if the weight to the disvalue of inequality is suitably chosen, egalitarianism claims that the levelling-down is always, all things considered, strictly worse. Furthermore, we can come up with an extensionally equivalent representation of egalitarianism that judges that the levelling-down is not better in any respect. From this, I would not take the Levelling Down Objection to be a serious threat to egalitarianism. As you can see, the example of egalitarianism above is a linear combination (more specifically, strictly S-concave function) of people's welfare. On the other hand, prioritarianism is committed to a strictly concave function of people's welfare: the moral importance of a person's welfare diminishes, as the absolute level of his welfare gets higher. It requires a measure for the absolute level of welfare. But I do not think that the normative notion of welfare is something that can be measured on an absolute scale. The real issue behind the debate between egalitarianism and prioritarianism seems to lie in whether the social welfare function should be linear or not, or strongly separable or not.

One additional note is that Holtug rightly distinguishes the Person-Affecting View from the Pareto principle. The Person-Affecting View (or Holtug's Narrow Person-Affecting Principle) holds that an outcome X cannot be better (or worse) than another outcome Y if there is no one for whom X is better (or worse) than Y. On the other hand, the (strong) Pareto principle holds that an outcome X is strictly better than another outcome Y if (not 'only if') X is better for some person than Y and X is worse for no person than Y. Obviously, the Person-Affecting View and the Pareto

principle are logically independent. This is an important point, but one that a surprisingly small number of philosophers recognize.

I have pointed out two relative disagreements with Holtug's arguments. But these two points do not undermine the comprehensive theory of justice that he puts forward. Holtug's *Persons, Interests, and Justice* is a magnificent book. It exhibits a most thorough analysis of the most difficult problems in contemporary moral philosophy. His vigorous project deserves the highest acclaim. Any moral philosopher, actual or possible, ought to examine this book carefully.

Iwao Hirose
McGill University

doi:10.1017/S0266267112000077

Rational Choice, Itzhak Gilboa, MIT Press, 2010, xv + 158 pages.

Itzhak Gilboa's book aims at introducing rational choice theory to a readership without prior knowledge of the field. It presents the fundamental ideas and concepts, not, or very little, the mathematics behind them. Gilboa gives priority to intuitive explanations and illustrative examples. The mathematical details are relegated to the online Appendix.

Rational Choice Theory encompasses decision theory, game theory and social choice theory, as practiced today not only in theoretical economics, but also in computer science, logic and philosophy. One of the main assets of the book is in fact to show how close the concerns in rational choice theory are to those in many areas of philosophy.

This is a very good book, but one that should be 'Rated PG', or rather LG for 'lecturer's guidance'. The highly pedagogical introduction to the key concepts of Rational Choice Theory touches a lot of fundamental questions in philosophy, but there are not enough references for the student or the newcomer to see the points of contact clearly. This review tries to bring some of them into light.

Individual Decision Making: The first part of the book touches issues classically pertaining to Decision Theory. Chapter 1 introduces the basic concepts, the most crucial one being, unsurprisingly, rationality. Gilboa's claim that rationality, as understood in rational choice theory, is also accepted by 'most psychologist and behavioral decision theorists', is refreshingly controversial, and to a large extent correct, although it doesn't quite do justice to the literature on 'ecological rationality' (cf. Gigerenzer and Selten 2002). His own view on rationality is subjective and dynamic: 'a mode of behavior is rational for a given person if [he] feels comfortable with it, and is not embarrassed by [read here: would not be willing to