

the wanderer suggested in Friedrich's original nor of his sublime view (Buchka 1983: 52–53). Perhaps my reading of Bronk's book is biased by the expectation that I would see something of this transformation instead of reading a book that tried to tie our contemporary world to a world forlorn, one of which already Goethe was aware was unlikely to come back.

Harro Maas
University of Amsterdam

REFERENCES

- Buchka, P. 1983. *Augen kann man nicht kaufen: Wim Wenders und seine Filme*. München: Hanser.
- Cannon, S. F. 1978. *Science in Culture: The Early Victorian Period*. New York: Science History Publications.
- Henderson, J. P. 1996. *Early Mathematical Economics: William Whewell and the British Case*. London: Rowman & Littlefield.
- Kehlmann, D. 2005. *Die Vermessung der Welt*. Berlin: Rowohlt.
- Todhunter, I. 1876/2001. *William Whewell: An Account of his Writings, with Selections from His Literary and Scientific Correspondence*. Bristol: Thoemmes.
- Whewell, W. 1837. Letter to Charles Babbage. *The Athenaeum*, May 30.

doi:10.1017/S0266267110000246

Theory of Decision under Uncertainty, Itzak Gilboa. Cambridge University Press, 2009. xiv + 215 pages.

Hundreds or even thousands of academics working in universities all around the world consider themselves to be decision theorists. However, most decision theorists do not work in departments that contain the word 'decision' or 'choice' in their names. Why is that? Why do so few universities have a 'Department of Decision Theory'? At present, most people working in the field are enrolled as economists, psychologists, philosophers, computer scientists, or statisticians with a special interest in decision making.

No matter what the reason for this might be, modern decision theory is a truly *multidisciplinary* subject. However, it is far from clear that it is also an *interdisciplinary* one. People coming from different disciplines work on roughly the same problems, but they do not seem to communicate or cooperate very much with each other. As a consequence, rather few books and papers are interdisciplinary in the literal meaning of that word. There is simply too little exchange going on across the disciplinary boundaries.

Itzak Gilboa's *Theory of Decision under Uncertainty* is written with the ambition to fill (parts of) this disciplinary gap. Gilboa, who is an

economist by training but is now working in one of the few academic departments with the word 'decision' in its name, explains in the preface that the book is based on lecture notes for his graduate class on decision making under uncertainty. Gilboa has a very ambitious view about what material is relevant for such a course. Apart from covering obvious topics such as the theories of von Neumann-Morgenstern, de Finetti, and Savage, the author also discusses free will, Goodman's grue-bleen paradox, Hempel's paradox of confirmation, and logical positivism. These are topics that are traditionally dealt with in philosophy of science courses, but Gilboa claims that insights from decision theory can help clarify them.

Unlike David M. Kreps' classic *Notes on the Theory of Choice*, Gilboa explicitly aims to cover both the technical and the non-technical aspects of decision theory. In fact, one of the major strengths of the book is that it summarizes ideas and results from economics, psychology, and philosophy. However, Gilboa's aim is not just to summarize the literature. He also seeks to put forward his own views on a number of important topics. In what follows I will comment on some of the claims defended by Gilboa. I will focus mainly on the claims on which I disagree.

Chapter 2 is largely a discussion of free will and determinism. Most decision theorists just take for granted that the decision makers have free will, but for Gilboa this is an open question that needs to be addressed (and many philosophers would of course agree). However, his discussion of free will is related to a discussion of determinism in a somewhat surprising way. Gilboa argues that, 'One may model the world as deterministic . . . [b]ut this model will be observationally equivalent to another model in which the world is nondeterministic. It follows that, as far as we will ever be able to know, the world is not deterministic' (p. 6). I find this point problematic. Because even if Gilboa's argument was valid, it would not support the claim that we (as far as we will ever be able to know) have free will. Nondeterminism does not imply free will. That some events are genuine random events does not mean that they are under my control. Randomness is not free will. Moreover, quite a few philosophers have defended a view called 'soft determinism', according to which free will is *compatible* with determinism.

About half of the book is spent discussing a number of preference-based views of (subjective) probability and the concepts of utility related to those views. Unsurprisingly, Savage's view plays a major role, and Gilboa's exposition is very clear and concise. The student who first reads Gilboa's summary will be very well prepared for tackling Savage's text. However, I found it a bit odd that Jeffrey's version of subjective expected utility theory is not mentioned at all, especially since it is more general and conceptually more precise than Savage's. Jeffrey is only mentioned in a footnote, in the discussion of Newcomb's problem.

Unlike most non-philosophers interested in decision theory, Gilboa takes Newcomb's problem seriously: you are offered a choice between two boxes. The first box contains \$1000 and you know this, because the box is transparent and you can actually see the money. The second box contains either a million dollars or nothing. This box is not transparent, so you cannot see its content. Now imagine a Predictor who is very good at predicting other people's choices. Ninety-nine per cent of all predictions made by the Predictor so far have been correct. You are invited to make a choice between two alternatives. You either take what is in both boxes, or take only what is in the second box. You are told that the Predictor will put \$1M in the second box if and only if she predicts that you will take just the second box; otherwise that box will be left empty. The Predictor first makes her prediction and then puts money into the boxes; thereafter you make your choice. What should you do?

The point about Newcomb's example is, as some readers may already know, that the dominant strategy – to take both boxes – seems to be in conflict with the principle of maximizing expected monetary value. If you take both boxes your subjective probability will of course be very high that you will get just \$1000. If you ignore the dominance principle and take just one box, you can expect to win a much larger amount. Unlike most scholars working on Newcomb's problem, Gilboa argues that a rational decision maker should take only the first box. His argument is based on the claim that there are four relevant states to consider:

- (1,1) – the money is there anyway;
- (1,0) – \$1M are found only by the greedy;
- (0,1) – \$1M are given to the modest alone;
- (0,0) – the non-transparent box is empty

The decision matrix will now, according to Gilboa, look as follows:

	(1,1)	(1,0)	(0,1)	(0,0)
Greedy	\$1 001 000	\$1 001 000	\$1000	\$1000
Modest	\$1 000 000	\$0	\$1 000 000	\$0

Gilboa claims that, 'Clearly, there is no dominance in this matrix. There exists a state, (0,1), in which modesty pays off' and given that the probability of that state is high, one should take one box. On Gilboa's view, this 'resolution is very compelling' (p. 115).

I strongly disagree with Gilboa on his solution to Newcomb's problem. I do not think that he has provided a compelling one-box resolution of Newcomb's problem. As far as I can see his argument is based on a trivial mistake. The mistake has to do with the way Gilboa sets up the decision matrix. Imagine that the Predictor has already made her prediction. The \$1M either is in the second box or not. You are just

about to make your choice, and your choice will of course not affect the likelihood that the second box is empty. Since this is an important decision, you decide to draw a decision matrix and think things through carefully. Gilboa advises you to include the state '\$1M are found only by the greedy' in your decision matrix (as well as the parallel state '\$1M are given to the modest alone'). Let us suppose for the sake of the argument that the former state, i.e. '\$1M are found only by the greedy', is the true state of the world. It then follows that *if* you decide to choose 'Greedy' (take both boxes), *then* the second box will not be empty. However, *if* you decide to choose 'Modest' (take just the first box), *then* the second box will be empty. This shows that according to Gilboa the actual content of the second box will *depend on what you decide*, contrary to what is specified in all interesting formulations of Newcomb's problem, including the original one given by Nozick:

The situation is as follows. First the being makes its prediction. Then it puts the \$1M in the second box, or does not, depending upon what it has predicted. Then you make your choice. What do you do? (Nozick 1969/1997: 46)

I believe this demonstrates that Gilboa's resolution is based on a misunderstanding of the way the problem is set up. A convincing defence of the nowadays unfashionable one-box solution (as suggested by evidential decision theory) must at least respect the basic ideas behind the formulation of the problem.

The chapter on frequentist definitions of probability (Chapter 4) discusses in considerable detail what reasons we have for thinking that empirical frequencies for past events might tell us anything about probabilities for future events. Are we really justified in thinking that the future will resemble the past? If not, frequentist definitions of probability seem to be false. As Gilboa correctly points out, this problem directly leads us to David Hume's riddle of induction, as well as to Nelson Goodman's new riddle of induction (the grue-bleen paradox). Gilboa devotes a major part of the chapter to a discussion of Goodman. The main point is that there is, according to Gilboa, a simple and convincing resolution of Goodman's grue-bleen paradox: Although the grue/bleen vocabulary is no more complex than the blue/green vocabulary, Goodman's way of setting up the problem does not allow us to acknowledge the fact that the terms blue and green, but not grue and bleen, 'have explicit reference, or [are] proper names' (p. 29) in English. However, as far as I can see, Gilboa's proposal does not solve the problem. Suppose, for instance, that we come across a tribe in the jungle who speak another language, Frenglish, in which the terms grue and bleen have exactly the same properties as blue and green in English. It now seems that we are back where we started.

My general impression is that the parts of the book most closely related to the author's field of expertise are of a very high quality, and well worth reading for graduate students as well as others interested in decision theory. However, some of the chapters that touch upon issues that lie outside the author's field of expertise, such as philosophy, are weaker, as exemplified above. When reading those chapters it might be helpful for the reader to also take a look at other books or articles. This does not mean, however, that I do not think that this is an important and valuable addition to the literature. Decision theory is a multidisciplinary subject, and all attempts to make it more interdisciplinary should be welcomed.

Martin Peterson

Eindhoven University of Technology

REFERENCE

Nozick, R. 1969/1997. Newcomb's problem and two principles of choice. In *Socratic Puzzles*. Cambridge, MA: Harvard University Press.

doi:10.1017/S0266267110000258

Liberty, Games and Contracts: Jan Narveson and the Defence of Libertarianism, Malcolm Murray (ed.). Ashgate, 2007. 273 pages.

Jan Narveson is well known for his defence of right-libertarianism on contractarian grounds (e.g. Narveson 1988). In this volume, a Festschrift in his honour, friends and students of Narveson critically evaluate Narveson's theory.

Narveson's position can be summed up in three fundamental claims. First, the justification of a political philosophy or indeed any normative ethical theory, requires contractarian foundations. All contractarians consider morality as the outcome of an agreement among relevant parties. More precisely, moral norms are those rules that are agreed upon by agents in a suitably characterized bargaining situation. Contractarians share this starting point with other social contract theorists. However, contractarians differ from other social contract theories, like that of John Rawls, in that the latter treat such an agreement among rational agents as a heuristic instrument for identifying the content of morality. That is, authors like Rawls claim that moral norms are binding for reasons other than that they are agreed upon by agents in the original position. Narveson, like other contractarians, believes that agreement of some sort is necessary and sufficient for the normativity of such norms. ('Of some sort' because closer reading reveals that this social contract is not an actual