speaking of instrumentalism is justified. Because of its fruitfulness, rational choice modelling may be justified along that line, but so may also be the use of models in which rational choice is qualified or even given no role at all.

The purpose of the foregoing digression is to suggest that, as 'nuts and bolts', a comfortable habitat for mainly psychological mechanisms might be located not only in the real world but also in the models that economists deem necessary to discuss some social phenomena. This would have been impossible two decades ago because all mainstream economists were still shy of departing from strict rational choice as the main feature of their constructs. This is not true anymore. In Bénabou and Tirole (2006), for instance, several of the psychological mechanisms presented by Elster – especially dissonance reduction and hyperbolic discounting (in the form of salience of the present) – play a major role to account for a number of complementary differences between the United States and Europe. Thus *Explaining Social Behavior* should be recommended also to young mathematical social scientists in search of inspiration.

Pierre Salmon
Université de Bourgogne

REFERENCE

Bénabou, R. and J. Tirole 2006. Belief in a just world and redistributive politics. *Quarterly Journal of Economics* 121(2): 699–746.

doi:10.1017/S0266267109990125

Preference and Information, Dan Egonsson. Ashgate, 2007, xi+163 pp.

This book provides a thorough and engaging treatment of one important consideration regarding the preferences of individuals – the standing of these preferences vis-à-vis the factual information upon which they are based. While the book is of wide interest, I agree with the claim of the author, Dan Egonsson, that it is particularly important for those who are concerned with 'preference-sensitive' measures of well-being or broader moral theories. Indeed, any moral theory that gives some importance to the preferences of individuals will need to engage with the issue of what sort of preferences count. It is common to appeal to the degree of autonomy of a person's preferences, but what exactly does this mean? Some have focused on preference formation, the idea being that we should not pay much attention to preferences that have been somehow coerced or conditioned (see, e.g., Elster 1982). In this book, Egonsson focuses on a different part

https://doi.org/10.1017/S0266267109990125 Published online by Cambridge University Press

of the equation – the relationship between autonomy and relevant factual information.

The initial question is: what are ideal preferences, from the information point of view, or in other words, what would it mean for preferences to satisfy the 'full information requirement'? Egonsson introduces a number of distinctions in dealing with this issue – quantity versus quality in belief, instrumental versus intrinsic desire, preferences that depend on their own persistence and those that do not. Arguably the crux of the book, however, comes later, when Egonsson considers the relationship between actual and ideally informed preferences. He asks whether a disparity between the two is always regrettable (other things being equal), and how any such disparity should affect our moral calculations. In what follows, I discuss Egonsson's treatment of these themes in turn.

The opening examples of the book (Chapter 1) do a good job in probing our intuitions about why it is important for preferences to be suitably informed. Egonsson appeals to cases of disappointment or regret – a person acts on a preference and then later is disappointed when their preference is satisfied. One of the examples Egonsson uses is that of someone wanting to be an angst-ridden philosopher (à la Wittgenstein) and then being disappointed when they actually become one. Disappointment or regret is a sign that the person's preferences would have been different in the first place, had they known more, or in other words, disappointment is a sign that the person's preferences fell short of the 'information requirement'. (Note that Egonsson in fact distinguishes several versions of the information requirement; the one I make reference to here is the one I consider most pertinent to the questions addressed later in the book.) Egonsson is careful, however, to point out that disappointment is only a guide to information shortcomings; it may well be the case that the person has the full stock of relevant information and even predicts their own later disappointment, yet nonetheless persists with their current preference. In line with Parfit, Egonsson refers to these as preferences that 'are not implicitly conditional on their own persistence' (p. 8). So the real test with respect to the information requirement is not the presence/absence of later disappointment, but rather whether the person's preferences would change if they were to know more at the time.

In subsequent chapters (2–4), Egonsson dwells on what 'knowing more' really amounts to. The straightforward view of what it means to be fully informed is to have all the relevant true beliefs and no false beliefs. Where Egonsson deviates from the standard line is his insistence (in Chapter 3) that it is not just the quantity but also the quality of our beliefs that matters. Here quality is not about the information content of a belief (that is taken care of in the 'quantity' side of things); rather, quality is supposedly about how well we internalize our beliefs, as opposed to just representing them propositionally. The idea is that there is variation in how

https://doi.org/10.1017/S0266267109990125 Published online by Cambridge University Press

well we appreciate or 'vividly imagine' the content of our beliefs. Egonsson suggests, for example, that we tend to better imagine momentary sensory experiences, as compared with more enduring emotional experiences. According to Egonsson, this is relevant to the information requirement because it is possible that a person's preferences would change if only they were to better internalize/imagine their own set of true beliefs.

Egonsson's appeal to the varying quality of our beliefs is interesting, but this seems an unnecessary complication in our model of belief. To begin with, there is reason to be cautious about any intuitions we might have regarding the varying prominence of our beliefs; it is plausible that any difference in intensity stems not from the belief side of things, but rather from differences in our *desiring attitudes* towards the content of our beliefs. Moreover, any case that seems to involve a deficiency in quality of belief can arguably be just as easily described in terms of a deficiency in quantity of belief. For instance, Egonsson has a very interesting discussion in Chapter 3 of the impossibility of properly imagining a time-extended event in any amount of time that is less than the time the event itself will endure. I like the spirit of this discussion, but I think it can just as well be put in terms of the quantity of true beliefs – it takes time to build up a suitably fine-grained stock of propositional beliefs regarding a possible event; indeed, perhaps this would take an ordinary person the same or even *more* time than the duration of the event itself. So Egonsson's overarching point here still stands – it may be impossible for anyone to actually construct fully informed preferences (even if their preferences would nonetheless withstand full information). Considerations of belief quantity alone tell us that attending to all the relevant details of possible future events may not be a realistic undertaking.

Egonsson considers the issue of what are the objects of preference. According to one view, labelled 'intentional', the objects of preference have just whatever properties the agent ascribes to them; such preferences would seem to be immune to facts about the world and thus the information requirement. For example, if John wants to be a philosopher who never does any administrative work then there is nothing wrong with John's original preference if it turns out that philosophers invariably have to do some administration on top of research. The 'extensional' view, by contrast, takes the objects of preference to be things in the world to which we either correctly or incorrectly ascribe properties. This is how we naturally think of objects of preference when discussing the information requirement. On this view, John's preference was misguided in that he thought being a philosopher does not involve administrative work when in fact this is not (generally) the case. Egonsson argues, however, that even if we understand preferences according to the intentional view, they can still fail the information requirement if we consider necessary consequences of the objects of our preferences to be bound up with the preferences

https://doi.org/10.1017/S0266267109990125 Published online by Cambridge University Press

themselves. In my opinion, this move simply blurs the distinction between the intentional and extensional views. Having said that, I do not find the intentional view of the objects of preference very convincing to begin with. There does not seem much use for a view in which preferences need not match up to the way the world is. Indeed, this is contrary to the decision-theoretic model in which preferences are always about actual consequences in the world, or acts that are associated with possible consequences (see Jeffrey 1965). Egonsson himself notes that the extensional view is surely the more fundamental (p. 12); his interest in the intentional view seems more for the purpose of showing that the information requirement is a problem for everyone, that is, whatever one thinks about the objects of preference.

There is another distinction concerning preferences that Egonsson makes quite a lot of: the difference between instrumental and intrinsic value. The traditional line is that valuing something instrumentally is to value it as a means to something else. Valuing something intrinsically, on the other hand, is to value it for 'what it is in itself'. Egonsson tells a more nuanced story than that. In Chapter 7 (pp. 79–80), he discusses Korsgaard's proposal that there are in fact two distinctions: finality and instrumentality on the one hand, and intrinsicality and extrinsicality on the other. For example, I might attribute final value to my grandmother's furniture (i.e. I do not wish to use it to achieve something else) and yet the furniture does not have intrinsic value given that the value is derived from its relation to my grandmother. Egonsson refers to a preference like this (for, say, grandmother's chair) as a conditional or weakly intrinsic preference. The point is that such preferences may be mistaken if they rest on false beliefs about whether the relevant conditions hold (whether the chair really did belong to my grandmother). It is surely a bit odd, however, to attach the instrumental/intrinsic/etc. labels to whole preferences, since objects of preference can instantiate a number of different values. For example, a chair might have instrumental value (in that it facilitates sitting still without getting a backache), final extrinsic value (in that it is an heirloom), and also final intrinsic value (in that it is a beautiful object). In such cases it is unclear what sort of preference one has for the chair. That said, what Egonsson seeks to illustrate is that the information requirement encompasses a variety of situations of preference (i.e. different sorts of preference would be overturned if the agent was better informed), and I think his examples show that this is indeed the case, even if one does not necessarily agree with the finer points of Egonsson's categorization of preferences.

So the practical question is: what is the significance of the information requirement? Why does it matter if an agent's preferences would be overturned if they were better informed? Does this make the agent and/or their preferences irrational? From the very beginning of the

https://doi.org/10.1017/S0266267109990125 Published online by Cambridge University Press

book (p. 9), Egonsson refers to preferences that do not satisfy the information requirement as being less than fully rational. This is in line with other writers, like Harsanyi. In my opinion, this an unfortunate use of terminology, because the preferences that an agent holds might be perfectly reasonable given their epistemic circumstances. As a deliberating agent, we never have access to all the facts, and so rational preferences (especially with regard to acts) are just those that maximize *expected utility* (rather than actual utility) relative to our beliefs and desires at the time. (I note that Egonsson has some qualms about expected utility theory. This is just a passing point in the book, however, and nothing rests on what is taken to be the correct theory of rational choice.) Egonsson does draw attention to the difference between subjective and objective standards when it comes to preference 'rationality', but this discussion is very late in the book (in Chapter 6). He refers to comments by Persson and Brandt (pp. 72–73) regarding the difference between what is rational for a person and what is good for them, or what is subjectively rational versus what is objectively rational. The idea is that we can assess preferences either from the point of view of an agent who is deliberating, or else from an outside objective standpoint. If the information requirement is to have any real impact then it will be with the latter kind of assessment. Indeed, gathering more information is just another option in the decision-theoretic setting that may or may not be rational for an agent to choose.

One issue that does seem relevant both to the measurement of well-being and to an agent's own deliberation, however, is whether it is always desirable for preferences to be fully informed (even assuming, in the deliberative setting, that the information is not costly to obtain). Egonsson draws attention (in Chapters 4 and 5) to examples that have the character of a 'lucky dip'. These are things we want precisely because we do not know all there is to know about them. Moreover, if we were to know all the details then our preferences might change in a way that results in less satisfaction or utility. For instance, Egonsson considers a preference for climbing Mount Everest that at least partly rests on the adventure involved; it would ruin the excitement of the activity to know what lay beyond every turn and what all the difficulties would be. Some other slightly different kinds of examples are also discussed (some attributed to Gibbard). Egonsson does not present any firm opinions about these cases; he rather outlines the different positions one might take. The important thing is that there are a number of cases that collectively put pressure on the view that ideal preferences are fully informed preferences.

There will be many unambiguous cases, however, in which it is clear that an agent's preferences would change for the better if only they were better informed. The problem of the final chapters of the book (9 and 10) is how this should affect our moral calculations. Egonsson's arguments in this section do not support the information requirement having any

https://doi.org/10.1017/S0266267109990125 Published online by Cambridge University Press

practical authority. He focuses on a particular sort of moral quandary – when an agent has the power to honour or else disregard the (less informed) preferences of another. The examples are constructed so that all other things are equal – it is not the case that the other person would be grateful afterwards for having their initial preferences frustrated. For instance, Egonsson speaks of a father who wants to give his daughter a CD as a present; she likes punk but he thinks she would prefer classical if only she had more musical training. We are told that if the daughter were given a classical CD, she would never listen to it and so it would have no positive impact on her life. Egonsson argues that the information requirement does not provide any good reason to honour this kind of preference – a merely hypothetical preference – for the classical CD. He considers accounts from Railton and Rosati that could be perceived as toned down versions of the information requirement, in that they tend to give results in line with our intuitions (that the girl should receive the punk CD). But again, Egonsson finds these accounts problematic because they still permit an agent's actual preferences to be overridden without any benefit to the agent.

It might be argued, however, that Egonsson's examples put the information requirement in a bad light. If we considered cases where a person would later be grateful for having their actual preferences overridden, then we might look upon the information requirement more favourably. Perhaps it is the two criteria together that determine when overriding preferences is acceptable or in the interest of autonomy: the preferences fall short of the information requirement *and* the person involved would be grateful later for the intervention. Of course some modifications to this account are necessary to deal with the important practical problems, such as euthanasia preferences, that Egonsson discusses in Chapter 10 – here there is clearly no possibility that the person would be grateful later for any kind of intervention. At this point, Egonsson introduces a further distinction – the difference between preference and consent.

Like a number of the problems raised in the book, Egonsson leaves the more practical issues in medical ethics rather open-ended. Indeed, he makes explicit in the conclusion of the book that his goal was to identify and discuss the spectrum of problems associated with the information requirement, rather than to pursue a particular line of argument. It is true that Egonsson does not take a strong position on any of the more subtle issues discussed in the book. Having said that, one could say that Egonsson leaves the reader with an overall bad impression of the information requirement (and for good reason!). Despite its potential relevance to a broad range of preferences, the requirement has restricted application – it concerns assessments of what preferences are objectively good for an agent to hold, rather than what preferences are rational for them to hold. Moreover, Egonsson suggests that it is not always good for an agent to

https://doi.org/10.1017/S0266267109990125 Published online by Cambridge University Press

satisfy the information requirement. Finally, he argues that the fact that an agent's preferences do not satisfy the information requirement does not, on its own, justify our disregarding these preferences in our moral calculations. This would detract from the agent's autonomy, rather than enhance it.

Katie Steele
University of Sydney

REFERENCES

Elster, J. 1982. Sour grapes – utilitarianism and the genesis of wants. In *Utilitarianism and Beyond*, ed. A. Sen and B. Williams, pp. 219–38. Cambridge: Cambridge University Press.
Jeffrey, R. 1965. *The Logic of Decision*. McGraw-Hill.

https://doi.org/10.1017/S0266267109990125 Published online by Cambridge University Press