remarkable book is important not only in virtue of the extent to which it already demonstrates that fecundity, but also in virtue of the further work on shared agency to which it is sure to give rise.

Andrea C. Westlund
University of Wisconsin-Milwaukee

http://dx.doi.org/10.1080/00048402.2015.1025800

Buchak, Lara, *Risk and Rationality*, Oxford: Oxford University Press, 2013, pp. xii + 256, £40 (hardback).

Lara Buchak has written a wonderfully original, clear, and compelling formal account of the role that risk might play in a normative theory of rational choice. Standard expected utility theory (SEU) represents a rational agent's preferences via two parameters: a probability function p and a utility function u. In Buchak's more liberal *risk-weighted expected utility theory* (REU), a third parameter r is introduced: a function that captures the agent's attitude toward risk. Whereas SEU mandates risk-neutrality, REU permits rational risk-aversion and risk-seeking.

The formal machinery for this account is adapted from two sources. One is Köbberling and Wakker [2003] henceforth 'KW'; the other is Machina and Schmeidler [1992] 'MS'. However, the idea of combining these two formal schemes is a highly original one, and indeed they dovetail surprisingly well. The resulting account is compelling in that it is clear from the formal details that the agent's attitude to risk has been neatly separated from her beliefs, so that Buchak's function p indeed deserves to be regarded as capturing degrees of belief. The clarity of exposition and the author's sensitivity to what is of philosophical importance make this book an exemplary piece of formal philosophy. (I will refer to Buchak's book below as 'B'.)

There are various ways in which this risk-weighted approach might be motivated. One way is via cases like Allais's famous example of apparently coherent yet risk-averse preferences [Allais 1953; B: 12]. Note, however, that Buchak does *not* want her scheme to permit the characteristic preferences elicited by the Ellsberg examples of 'ambiguity aversion' [Ellsberg 1961; B: 80]. This means finding a principled way of separating the Allais and Ellsberg cases. That will be key to what follows.

Of these various ways to motivate the risk-weighted approach, the simplest is, perhaps, this. Suppose that you and I share a preference for receiving one pie for sure, over a gamble g that yields two pies if the toss of a fair coin lands heads and nothing if it lands tails. But suppose that the explanations for our shared preference differ. I am full after a single pie and I have diminishing marginal utility for a second; you are insatiable but risk averse. SEU has no way of distinguishing these two distinct possibilities.

The formal set-up is a Savage-style decision theory with a tripartite scheme of *acts*, *states*, and *outcomes*. An act (or *option*) f is construed as a function from states to outcomes. The agent has preferences over the set of acts. We write $f<g$ when g is preferred to f. By construing outcomes as constant acts, we may extend the preference relation $<$ to outcomes. A set of states is called an *event*. Acts are assumed to be *simple*, in the sense that each yields only finitely many outcomes. Thus each act induces a partition $\{E_1,\ldots E_n\}$ of the state space into events, each of which is the inverse image under f of some particular outcome.

The idea of Savage's theory is to show that if the agent's preferences satisfy certain axioms then they may be given an expected utility representation with respect to a unique probability function p on the set of states, and a utility function u on the set of outcomes that is unique up to positive linear transformation. In a simple case like the pie example, in which there are two possible states of the world s and t, and two possible outcomes x and y, the standard expression for the expected utility of the act $f =$ [s,x; t,y] that delivers outcome x in state s and outcome y in state t, is this:

$$\mathrm{SEU}(f) = p(s).u(x) + p(t).u(y)$$

Now, if y is preferred to x, this can be rewritten in cumulative or stepwise fashion as the following:

$$\mathrm{SEU}(f) = u(x) + p(t)[u(y) - u(x)]$$

In the stepwise representation, the utility of an option is expressed as the baseline utility of receiving the worse outcome, plus the utility difference between the better and worse outcomes multiplied by the probability that one will do better.

In Buchak's system, the risk function r is then applied to that probability, yielding this:

$$\mathrm{REU}(f) = u(x) + r(p(t))[u(y) - u(x)]$$

In the pie problem, the two relevant states are *heads* and *tails*, and the outcomes, ranked for both of us in order from worst to best are *no pie*, *one pie*, *two pies*. The distinction between your risk-aversion and insatiability and my own risk-neutrality and diminishing marginal utility for pie can be captured, for example, by the assumption that my utility function for pie is $u(n) = \frac{1}{2}\sqrt{(2n)}$, where n is the number of pies received, while yours is $u^{*}(n) = n$. Suppose, in addition, that my risk function is the (risk-neutral) identity function $r(x) = x$, and yours is $r^{*}(x) = x^{2}$. The latter function expresses a risk-averse attitude in virtue of having the formal property of being concave upward. Note that for both of us $p(heads) = p(tails) = \frac{1}{2}$, since we agree that the coin is fair. Then, plugging these values into the REU formula, we find that, for me,

$$\mathrm{REU}(g) = 0 + r(1/2)[1 - 0] = 1/2 < 1 = \mathrm{REU}(one\,pie)$$

while for you,

$$\mathrm{REU}^{*}(g) = 0 + r^{*}(1/2)[2 - 0] = 1/2 < 1 = \mathrm{REU}{*}(one\,pie).$$

So our preferences are the same, though the explanations of those preferences differ.

One half of the technical machinery Buchak appeals to comes from KW's development of cumulative expected utility theory (CEU). The central notion is that of *tradeoff equality*, a four-place relation $xy{\sim}^{*}zw$ between outcomes that obtains when receiving x instead of y is just as good as receiving z instead of w. Operationally, this means that there are acts f and g, and some non-null event E (i.e. an event that will receive non-zero probability), such that the agent is indifferent between the options [x, E; f, $\sim E$] and [y, E; g, $\sim E$] and also indifferent between [z, E; f, $\sim E$] and [w, E; g, $\sim E$]. This is taken to be the criterion for $u(x)$-$u(y) = u(z)$-$u(w)$. To obtain a representation theorem for SEU via this method, an axiom of Tradeoff Consistency must be imposed. This axiom says that 'improving an outcome in any $\sim^{*}$ relationship breaks that relationship.' In other words, $xy\sim^{*}zw$ and $y < y'$ entails that it's not the case that $xy'\sim^{*}zw$.

To adapt this idea for a stepwise utility theory, KW relaxes the Consistency axiom. Two options *f* and *g* are said to be *comonotonic* when there is no pair of states *s* and *t* such that the outcome $f(s)$ is preferred to $f(t)$ and yet $g(t)$ is preferred to $g(s)$. A *comoncone* is a set of options, all of which are comonotonic. Then the *comonotonic tradeoff equality* $xy\sim_{C}^{*}zw$ obtains when '*x* and *y* play the same compensatory role as *z* and *w* in some gambles *f* and *g* where all of the modified gambles after *x*, *y*, *z*, and *w* have been substituted in are in the same comoncone' [B: 90]. The restricted axiom of Comonotonic Tradeoff Consistency (CTC) states that improving an outcome in any $\sim_{C}^{*}$ relationship breaks that relationship.

The central result of KW is a theorem [KW Thm 8: 399; B: 239] stating that when the agent's preferences satisfy a certain subset of Buchak's axioms, including, most critically, CTC, there exist a unique utility function *u* and a weighting function *w* that allow those preferences to be represented by a cumulative expected utility function. In the simple case where the state space is partitioned into two events *D* and *E*, and $f(D)<f(E)$, the expression for CEU(*f*) is as follows:

$$\mathrm{CEU}(f) = u(D) + w(E)[u(E) - u(D)]$$

The similarity with the expressions given above for SEU(*f*) and REU(*f*) should be obvious.

At the heart of the MS result is a familiar operational method for eliciting comparative probability judgments. Given a pair of events *D* and *E*, we want to know which one the agent considers more probable. We choose outcomes *x* and *y* such that $x<y$. Then we ask the agent to rank the following options: (1) [*x* if *D*; *y* if $\sim D$] (2) [*y* if *E*; *x* if $\sim E$]. Suppose the agent answers that she prefers (1) to (2). Then we may infer that she considers *E* more probable than *D*, for she would rather stake the better outcome *y* on the truth of *E* than on the truth of *D*.

For this method to be successful, we need a guarantee that the same probability comparison would have be elicited no matter what test outcomes had been selected in place of *x* and *y*. That guarantee is precisely the content of Savage's axiom P4, referred to by MS as the Weak Comparative Probability axiom (WCP).

But the imposition of WCP alone does not suffice to ensure that the comparative judgments so elicited are judgments about genuine probabilities. For example, the Ellsberg preferences satisfy WCP yet admit no probabilistic representation. Savage added to WCP the famous Sure-Thing Principle (STP); but, for MS and Buchak, that recourse is unavailable. That is because STP 'imposes consistency on beliefs *and* risk preferences' [MS: 762]. What is required, rather, is an axiom that imposes only consistency on beliefs, without implying full SEU. That is the precisely the role played by the axiom of Strong Comparative Probability, which MS use to replace Savage's STP and WCP.

SCP: If *D* and *E* are disjoint events, $x<x'$, $y<y'$ outcomes, and *f* and *g* acts, then if we elicit the information that $p(D)<p(E)$ by asking the agent to rank the options

[*x'* if *D*; *x* if *E*; *f* otherwise] and [*x* if *D*; *x'* if *E*; *f* otherwise],

exactly the same comparative probability judgment will be elicited by replacing *x*, *x'*, *f* with *y*, *y'* and *g*.

Where STP is violated by both the Allais and the Ellsberg preferences, SCP is violated only by the latter [MS: 763–4]. Thus, SCP serves perfectly to divide Allais from Ellsberg, that is, to permit rational risk aversion while still ruling out Ellsberg-style 'ambiguity aversion'.

It is no exaggeration to say that, from the formal point of view, it is SCP that is the key to Buchak's project. It can be proved [B: 241; MS Thm 8: 766] that any preferences that satisfy another subset of Buchak's axioms, including, most critically SCP, may be represented by a unique probability function p that has the nice property that the agent is indifferent between any two assignments of outcomes to states that give rise to the same probability distribution over outcomes, i.e.:

If $p(f^{-1}(x)) = p(g^{-1}(x))$ for all outcomes x, then f and g are equipreferred.

This result underwrites Buchak's definition of the risk function r. For each c $(0<c<1)$, let

$r(c) = w(E)$ for any E such that $p(E) = c$.

It is MS Thm 8 that establishes r as well-defined, since that is what allows us to show that any two events that have the same probability have the same weight [B: 242]. Finally, substituting r, so defined, into the KW definition of CEU, we obtain Buchak's REU. This completes the sketch of Buchak's representation theorem [B: 99, 238].

It is an impressive achievement. It is a compelling result, precisely because, as should be clear from the discussion above, the probabilities p elicited via the MS method really do represent genuine degrees of belief. Again, the secret to ensuring this neat separation of belief from risk-attitude is the appeal to the SCP, rather than to the conjunction of STP and WCP as in Savage's derivation of standard expected utility. The reader is also referred to Buchak's discussion of this key point in Section 3.4 [110–13].

This book opens up a discussion. Is Buchak correct in thinking that the Allais and Ellsberg cases should be treated differently, or should we prefer a more unified account like Isaac Levi's [1986], in which the Allais example is accommodated by allowing indeterminate utilities, and Ellsberg via indeterminate probabilities? Should we ask, as Brad Armendt [2014] has done, whether an account of risk should be *more* liberal than Buchak's by not tying the agent to a single risk function? These are questions worth thinking about. It is a virtue of Buchak's book that it lays the issues out so transparently and enables this discussion to be so clearly focused. Regardless of whether one agrees with Lara Buchak's proposal, she has provided us with an account of the role of risk to be reckoned with.

References

Allais, Maurice 1979 (1953). The Foundations of a Positive Theory of Choice Involving Risk and a Criticism of the Postulates and Axioms of the American School, in *Expected Utility Hypothesis and the Allais Paradox*, ed. Maurice Allais and Ole Hagen, Dordrecht: Reidel: 27–145.

Armendt, Brad 2014. On Risk and Rationality, *Erkenntnis* 79/6: 1119–27.

Ellsberg, Daniel 1961. Risk, Ambiguity, and the Savage Axioms, *Quarterly Journal of Economics* 75/4: 643–69.

Köbberling, Veronika and Peter P. Wakker 2003. Preference Foundations for Non-Expected Utility: A Generalized and Simplified Technique, *Mathematics of Operations Research* 28/3: 395–423.

Levi, Isaac 1986. The Paradoxes of Allais and Ellsberg, *Economics and Philosophy* 2/1: 23–53.

Machina, Mark J. and David Schmeidler 1992. A More Robust Definition of Subjective Probability, *Econometrica* 60/4: 745–80.

Savage, Leonard J. 1954. *The Foundations of Statistics*. Dover: John Wiley and Sons.

John Collins
Columbia University

http://dx.doi.org/10.1080/00048402.2015.1029500