*Review of Stefan Riedener, Uncertain Values: An Axiomatic Approach to Axiological Uncertainty**

Johan E. Gustafsson

Stefan Riedener, *Uncertain Values: An Axiomatic Approach to Axiological Uncertainty*, De Gruyter, 2021, 167pp, $16.99, ISBN 978-3-11-073957-2.

Stefan Riedener's book is concerned with axiological uncertainty — that is, the problem of how to evaluate prospects given uncertainty about what is the correct axiology. For evaluations of this kind of meta value, Riedener uses the term '*m*-value' (3). The main goal of the book is to provide an axiomatic argument for Expected Value Maximization, which is the view that an option x has an at least as great *m*-value as an option y if and only if x has an at least as great expected value as y, where the expected value of an option is a sum of the value of the option on each axiology weighted by one's credence in the axiology (5).

Riedener's argument (ch. 2-3) takes the form of a representation theorem, modelled after Harsanyi's social-aggregation theorem. The argument assumes that all axiologies and the *m*-value ranking satisfy the expected-utility axioms and that the following dominance condition holds (56):

> *Pareto Condition* For any two prospects with the same underlying probability distribution over axiologies, it holds that, if the prospects are equally good on all theories with non-zero probability, then the prospects are equal in *m*-value, and, if one of them is at least as good as the other on all theories with non-zero probability and strictly better on some, then it has a greater *m*-value.

Given these assumptions, Riedener proves a representation theorem showing that *m*-value must conform to Expected Value Maximization.

* Published in *Notre Dame Philosophical Reviews*, https://ndpr.nd.edu/reviews/uncertain-values-an-axiomatic-approach-to-axiological-uncertainty.

Riedener also provides a constructivist account of intertheoretic comparisons of value (ch. 4), extends the representation theorem to provide a formal separation of credences and values (ch. 5), and provides a variant of the theorem that relaxes the completeness assumption, which is one of the expected-utility axioms (ch. 6).

Riedener's book provides the best exposition of Expected Value Maximization to date. Still, I wish to bring up two major worries about the view — namely, the problem of higher-order inconsistency and the problem of how to make non-arbitrary intertheoretic comparisons of value.

We start with the problem of higher-order inconsistency. Suppose that you have a 2/3 credence in the axiological theory T_1 and a 1/3 credence in the axiological theory T_2. And suppose that you have a 2/3 credence in Expected Value Maximization and a 1/3 credence in My Favourite Theory (that is, the rival view that the m-value of an option is equal to its value according to the axiological theory in which you have the most credence). Now, consider a choice between two options a and b — whose values, according to T_1 and T_2, are as follows:

	T_1 (2/3)	T_2 (1/3)
a	5	14
b	10	1

According to Expected Value Maximization, the m-value of a is equal to $2/3 \cdot 5 + 1/3 \cdot 14 = 8$ and the m-value of b is equal to $2/3 \cdot 10 + 1/3 \cdot 1 = 7$. And, according to My Favourite Theory, the m-value of a is equal to 5 and the m-value of b is equal to 10. So, to maximize m-value according to Expected Value Maximization, you must do a, but, to maximize m-value according to My Favourite Theory, you must do b. Although you have higher credence in Expected Value Maximization than in My Favourite Theory, you are uncertain whether the former is the right approach to axiological uncertainty. So, to handle this meta uncertainty, we apply Expected Value Maximization to the uncertainty about the m-value of a and b. Riedener uses the term 'm^2-value' for evaluations of this kind of second-order axiological uncertainty (10). According to Expected Value Maximization applied to this meta level, the m^2-value of a is equal to $2/3 \cdot 8 + 1/3 \cdot 5 = 7$ and the m^2-value of b is equal to $2/3 \cdot 7 + 1/3 \cdot 10 = 8$. So, to maximize m^2-value according to Expected Value Maximization, you must do b. But, now you're in a bind, since you cannot maximize both m-value and m^2-value.

As long as you have some credence in Expected Value Maximization and that credence falls short of certainty, there will be possible cases like this. But, as Riedener admits, it seems 'almost indisputable' that we're not certain about Expected Value Maximization (54). So it seems that Expected Value Maximization cannot be a general solution to axiological uncertainty.

Riedener addresses the related issue of the regress problem — that is, whether there is an answer to what option has the greatest overall-value given uncertainty about m-value, m^2-value, m^3-value, and so on (10). (His response is that there is no such overall-value.) But the worry here, the problem of higher-order inconsistency, is that the agent can't maximize both m-value and m^2-value.

Note that many rival approaches to axiological uncertainty avoid this problem. As long as you have more credence in My Favourite Theory than in any other approach to axiological uncertainty, the acts that have the greatest m-value will also have the greatest m^2-value and the greatest m^3-value and so on. And, given externalism (that is, the view that the m-value of an option is equal to its value on the true axiology), the best option according to the true axiological view will have the greatest m-value, m^2-value, m^3-value, and so on.

It may be objected, that we can avoid this problem if we aim to maximize m^i-value only at those levels i such that the options with the greatest m^i-value also have the greatest m^j-value for all levels $j > i$. This view, however, is markedly different from Expected Value Maximization. And, as noted by Riedener (55), the Harsanyi-style defence is implausible for higher-level axiological uncertainty. The trouble is that the interesting rivals to Expected Value Maximization violate some of the expected-utility axioms or the Pareto Condition. And then the argument is unsound.

It may next be objected that we could adopt Expected Value Maximization for m-value while we adopt some other approach for m^2-value (for instance, externalism). But, if that other approach is appropriate for m^2-value, why wouldn't it be appropriate for m-value too?

Next, we turn to the problem of how to make non-arbitrary intertheoretic comparisons of value. Suppose that c is an outcome with a billion people with a well-being of 3 and that d is an outcome with the same people as c still with a well-being of 3 and, in addition, a billion people with a well-being of 1. And suppose that your credence is split between Average and Total Utilitarianism. According to Average Utilitarianism, the value of an outcome is equal to its average well-being, so the value of c is 3 and the value of d is 2. According to Total Utilitarianism, the value of an outcome is equal to its sum total of well-being, so the value of c is 3 billion and the value of d is 4 billion.

Expected Value Maximization requires that we compare the difference in value between these two outcomes according to Average Utilitarianism with the difference in value between the outcomes according to Total Utilitarianism. But, at first blush, it's far from clear how this could be done in a non-arbitrary manner.

Riedener suggests a number of norms governing credence distributions for intertheoretic comparisons of value (72–3). The first is Simplicity:

> *Simplicity* Other things being equal, you should favour simpler credence distributions over more complex ones.

It's not entirely clear how this principle is meant to help us. Riedener explains, using the example of a comparison between a theory where only pleasure is valuable and a theory where both pleasure and beauty are valuable:

> It's difficult to spell out precisely what 'simple' means. But we arguably have an intuitive understanding of it. So suppose Simplicity holds. Then it constrains the intertheoretic comparisons you can reasonably make. The credence distribution on which the value of pleasure is equally great on the pleasure- and the pleasure/beauty-theory is arguably simpler than that on which their ratio is 113.27, or anything other than 1. (73)

Applying this suggestion to our example, it may seem that we should favour the credence distribution where the ratio between the units of the two theories is 1 (that is, we take 1 unit of average well-being to be equal to 1 unit of total well-being). So, following Simplicity, it may seem that we should favour the comparison that the absolute difference in value between c and d on Total Utilitarianism is a billion times as large as the absolute difference in value between c and d on Average Utilitarianism.

So, unless your credence in Average Utilitarianism is a billion times higher than your credence in Total Utilitarianism, d will have a greater m-value than c. But to give so much more weight to the difference in value on Total Utilitarianism just because it is expressed in greater numbers seems arbitrary. After all, the additional people in d would reduce the average well-being by a third — which seems fairly significant by average-utilitarian standards.

This suggested comparison is the one that makes the value of outcomes on the two theories the same given a population of one person. It may be objected that we could get a more plausible result if we make the value of outcomes the same given some larger population. But, if so, what size should we use? Using the size of the current population won't work, since it will lead to dynamic inconsistency when the size of the population changes. Using the actual population won't work, since what population is actual may depend on what we choose — which is always the case in choices where the differences between Average and Total Utilitarianism matter.

The problem with arbitrariness is still worse for theories that do not overlap in terms of what kinds of things they value in outcomes. Suppose, for instance, that your credence is evenly split between Preference Utilitarianism (PU), which values outcomes by their sum total of preference-based utility, and Hedonistic Utilitarianism (HU), which values outcomes by their sum total of pleasure. And consider a choice between options *e* and *f*, which are valued as follows by the two theories:

	PU (1/2)	HU (1/2)
e	3	1
f	1	2

To follow Simplicity, it seems that we should favour the credence distribution where the ratio between the units of PU and HU is 1. Then the absolute difference between *e* and *f* is greater on PU than on HU, so Expected Value Maximization entails that *e* has a greater m-value than *f*.

The trouble is that the choice of unit for the scale of preference-based utility and the choice of unit for the scale of pleasure are both arbitrary. We may just as well choose a scale for pleasure which is like the first except that all well-being levels have been multiplied by 3. Let HU* be HU with this alternative scale. Then *a* and *b* are valued as follows by PU and HU*:

	PU (1/2)	HU* (1/2)
e	3	3
f	1	6

Given this arbitrary change in the unit of pleasure, the stakes are suddenly greater on HU* than on PU. So, following Simplicity, we again favour a ratio of 1 for the intertheoretic comparison. And then Expected Value Maximization entails that f has a greater m-value than e. Neither choice of unit for the scale of pleasure seems simpler or less arbitrary than the other. So an arbitrary change in the unit of the scale of pleasure led to an arbitrary change in what option had the greatest m-value. Hence it seems that Simplicity does not help us make intertheoretic comparisons of value between these theories in a non-arbitrary manner.

Riedener also suggests two other norms:

> *Conservatism* If you encounter new evidence, then of the possible changes to your credences that accommodate this evidence you should, other things being equal, favour less radical distributions over more radical ones.
>
> *Coherence* Other things being equal, you should favour more coherent credence distributions over less coherent ones. (73)

But it's doubtful whether Conservatism helps us — because, if you update arbitrary credences in a conservative way, then it seems that the resulting will be arbitrary too.

One reading of Conservatism is that it requires that you favour comparisons where you make smaller changes to your valuations of outcomes when you change your credence between axiological theories. Could this suggestion help us in the case of Preference and Hedonistic Utilitarianism? Suppose that you started off being convinced of Preference Utilitarianism and then came to have some credence in Hedonistic Utilitarianism. It may then seem unmotivated to suddenly think that everyone is much worse off (or much better off) than you previously thought. But the extent to which people prefer what gives them pleasure may differ from time to time and depending on what we choose. So we get the same problem as before with dynamic inconsistency or with intertheoretic comparisons of value depending on what we choose.

Likewise, Coherence does not seem to help us, since it's equally coherent to favour many different credence distributions for intertheoretic comparisons of value between Preference and Hedonistic Utilitarianism.

Moreover, Coherence seems to conflict with Simplicity. Let Beauty-over-Pleasure be an axiology that values both beauty and pleasure, with 1 unit of beauty being twice as valuable as 1 unit of pleasure. Let Only-Pleasure be an axiology that only values pleasure, and let Only-Beauty be an axiology that only values beauty. Finally, let Both-Equally be an axiology that values both beauty and pleasure, with 1 unit of beauty being equally valuable as 1 unit of pleasure. Given these axiologies, we cannot coherently believe (as suggested by Riedener's explication of Simplicity) that, if two axiologies overlap in the sense that there is just one kind of thing that they both value in outcomes, then the value of that kind of thing is equal on both axiologies. By Simplicity, it seems that pleasure on Beauty-over-Pleasure is as valuable as pleasure on Only-Pleasure, which is as valuable as pleasure on Both-Equally. On Both-Equally, pleasure is as valuable as beauty. And, by Simplicity, beauty on Both-Equal is as valuable as beauty on Only-Beauty, which is as valuable as beauty on Beauty-over-Pleasure. It follows that, on Beauty-over-Pleasure, pleasure is as valuable as beauty. But, as we specified that axiology, beauty is not as valuable as pleasure on Beauty-over-Pleasure.

Of course, there are other proposals for how to make non-arbitrary intertheoretic comparisons of value. One of the highlights of Riedener's book, however, is his rebuttal of those rival proposals (58–69).

In conclusion, while Riedener's book may not be fully successful in its defence of Expected Value Maximization, it does provide the best account to date of what this approach amounts to and what its commitments are. And, as mentioned, the discussion of rival accounts of intertheoretical comparisons of value is a great contribution to the literature.

I wish to thank Krister Bykvist, Stefan Riedener, Dean Spears, and Christian Tarsney for valuable comments.