Lecture 4: Choosing from a point of view

Richard Pettigrew *21st May 2025*

Handout (with references) & appendix: RICHARDPETTIGREW.COM/LOCKE

This lecture is based on joint work with Catrin Campbell-Moore and Jason Konek.

What came before...

(i) Norms for credences; (ii) teleological justifications for them; (iii) norms for inquiry.

...and what's to come

(iv) Norms for action.

Diachronic exploitability

Paolini is about to play a tennis match.

- You're approached first by a bookie offering you a £3 bet on her *winning* at a price of £1 (Decision 1); and
- You're approached by a bookie offering you a £3 bet on her *losing* at a price of £1 (Decision 2).[1]

[1] A £t bet on X pays £t if X is true and £0 if X is false.

Decision 1	w_1 *P wins*	w_2 *P loses*
Gamble 1	2	−1
Decline 1	0	0

Decision 2	w_1 *P wins*	w_2 *P loses*
Gamble 2	−1	2
Decline 2	0	0

Decline 1 + Decline 2 < Accept 1 + Accept 2.

Normative decision theories for different representations of uncertainty

Precise credences.

- *Savage-style expected utility theory (EU).*[2] No credences *permit* Decline 1 + Decline 2.
- *Risk-weighted expected utility theory (REU);*[3] *weighted-linear utility theory (WLU).*[4]. There are credences and risk attitudes that *demand* Decline 1 + Decline 2.

[2] (Savage, 1954).

[3] (Quiggin, 1982; Buchak, 2013)

[4] (Chew, 1983; Bottomley & Williamson, 2024).

Imprecise credences.[5]

[5] (Levi, 1974; Walley, 1991; Seidenfeld, 2004; Moss, 2015; Bradley, 2016).

- Γ-*Maximin*. There are credal sets that *demand* Decline 1 + Decline 2.
- *E-Admissibility*; *Maximality*. There are credal sets that *permit* Decline 1 + Decline 2.

No credences.[6]

[6] (Wald, 1945; Milnor, 1954; Hurwicz, 1951; Pettigrew, 2022).

- *Maximin*. *Demands* Decline 1 + Decline 2.
- *Hurwicz criterion*; *Generalized Hurwicz Criterion*; *risk-bounded utility theory*. There are risk attitudes that *demand* Decline 1 + Decline 2.

More generally, for any representations of uncertainty and attitudes to risk that don't result in expected utility, there will be sequences like the one above.

How bad is diachronic exploitability?

Possible defences and responses:

- *No Commodity.* It isn't always possible to find a commodity in which the bets can be made—our utility must be a linear function of quantities of it.[7]
- *Time-Slice Rationality.* There is nothing irrational about a guaranteed sure loss across different selves or time slices.[8]
- *Permissible vs Mandatory.* For E-Ad and Max: permissible sure loss, unlike mandatory sure loss, is no indication of irrationality.[9]
- *Resolute Choice/Intemporal Coordination.* We use our decision theory not to choose at each time, but to choose a strategy for picking at each time; rationality then demands we stick with that strategy at later times, even when we then would prefer not to.[10]
- *'By their fruits' evaluation.* The situations in which these decision theories give rise to sure loss are very specific, and rather rare. Perhaps the decision theories compensate for poor performance in those rare cases by their performance elsewhere.

[7] (Schick, 1986).

[8] (Christensen, 1996).

[9] (Seidenfeld, 2004).

[10] (Machina, 1989; McClennen, 1990; Gauthier, 1997; Buchak, 2013).

Sequence A
(50% probability you'll face this)

	P wins	*P loses*
	50%	50%
Decision 1A		
Gamble 1A	2	−1
Decline 1A	0	0
Decision 2A		
Gamble 2A	−1	2
Decline 2A	0	0

Sequence B
(50% probability you'll face this)

	P wins	*P loses*
	50%	50%
Decision 1B		
Gamble 1B	2	−1
Decline 1B	0	0
Decision 2B		
Gamble 2B	2	−1
Decline 2B	0	0

Suppose you are a Buchakian agent with the very risk-averse risk function $r(p) = p^3$. Then:

(i) You will decline the gamble in each first-order decision (i.e., you'll choose Decline over Gamble in 1A, 2A, 1B, 2B).

(ii) What's more, faced with the second-order decision between different implementable combinations of accepting and declining, you'll prefer your combination (i.e., always decline) over all the others.

So, even though your decision theory leads you to be exploitable via Sequence A, its performance in the face of Sequence B redeems it and makes the price of the exploitation worth it, from your point of view.

Self-Recommending and Self-Undermining Decision Theories

In general, if we are uncertain which decision problem we might face, we might formulate a higher-order decision problem.

At the first-order level, we have:

- *first-order states*, e.g., Paolini wins; Paolini loses, etc.
- *first-order options*, e.g., pay £1 for a £3 bet on Paolini winning; pay £7 for a £10 bet; decline to pay £1 for a £3 bet; etc.
- *first-order utility function*, e.g., the utility of paying £1 for a £3 bet on Paolini winning, if she wins; the utility of paying £7 for a £10 bet on Paolini winning, if she loses; etc.
- *first-order decision problems*, e.g., pay £5 for £10 bet on Paolini winning vs decline to pay that.

- *first-order picking strategies* a picking strategy takes a set of options and returns a probability distribution over them, e.g., faced with Decision 1, it might say 70% chance of picking Gamble 1 and 30% chance of picking Decline 1; faced with Decision 2, it might say 100% chance of picking Gamble 2 and 0% chance of picking Decline 2.

  (i) a picking strategy *picks for* a decision theory if it assigns positive probability only to options the decision theory considers permissible;

  (ii) a picking strategy *picks regularly for* a decision theory if it assigns positive probability to *all* the options the decision theory considers permissible.

At the second-order level, we have:

- *second-order states*, e.g. Paolini wins and I face the decision whether or not to pay £5 for a £10 bet on her winning; Paolini loses and I face the decision whether or not to pay £7 for a £10 bet on her winning, etc.
- *second-order options*, i.e., picking strategies.
- *second-order utility* of a picking strategy is the utility of the option it picks.

Now, suppose we are uncertain which decision problem we'll face. A particular probability distribution represents our uncertainty. Then we can ask a decision theory which picking strategies are permissible from its point of view. We say:

(i) A decision theory is *strongly self-recommending* if *all* the picking strategies that pick for it are *permissible* from its point of view.

  That is, if you use the decision theory, you'd be happy for anyone else who uses it to choose on your behalf.

(ii) A decision theory is *weakly self-undermining* if *some* of the picking strategies that pick for it are *impermissible* from its point of view.

  That is, if you use the decision theory, there are others who use it too whom you wouldn't be happy to choose on your behalf.

(iii) A decision theory is *strongly self-undermining* if *all* of the picking strategies that pick for it are *impermissible* from its point of view.

  That is, if you use the decision theory, you wouldn't be happy for anyone who uses it to choose on your behalf.

An example

You will face the choice whether to pay £t for a £10 bet on Paolini winning. That is, you will face the following decision problem for some unknown value of t:

	P wins	*P loses*
Accept	$10 - t$	$-t$
Reject	0	0

However, you don't know what the value of t will be. Indeed, let's say you have the uniform distribution over possible values of t. Then:

- Whatever your credence p that Paolini will win, EU with p will think EU with p is a rationally permissible way to choose; and indeed, it will think it is the only rationally permissible one.[11]
- Whatever your credence p that Paolini will win, REU with p and any risk function r other than the risk-neutral one will think it is better to use REU with a different credence or different risk function or both.[12]
- Whatever range of credences $[p, q]$ that Paolini will win, all probabilities between p and q will prefer to use EU with $\frac{p+q}{2}$ over a uniform regular picking function for E-Admissibility or Maximality.

[11] That is, if Verity maximizes expected utility and has credence 1/3 that Paolini will win, then she'll be happy to delegate her decision only to people who maximize expected utility with credence 1/3.

[12] That is, if Verity maximizes risk-weighted expected utility and has credence 1/3 that Paolini will win and risk function $r(p) = p^2$, then she won't be happy to delegate her decision to anyone who maximizes risk-weighted expected utility with credence 1/3 and risk function $r(p) = p^2$.

Some results

(i) EU is strongly self-recommending.

(ii) REU and WLU are sometimes strongly self-undermining.

(iii) E-Admissibility, Γ-Maximin, and Maximality are sometimes weakly self-undermining.

(iv) Maximin, Hurwicz Criterion, and the Generalized Hurwicz Criterion are sometimes strongly self-recommending.

What's so bad about being self-undermining?

Note: the objection from self-undermining is not vulnerable to the No Commodity, Time-Slice Rationality, or Permissible vs Mandatory responses. Moreover, it raises a criticism from a standpoint that is legitimate from the point of view of the theory being criticised—an internal critique.

Two worries about self-undermining theories:

- They create dilemmas: faced with a decision problem, should you do what the theory demands, or what one of the picking strategies it would have endorsed demands?
- They cannot be correct: if the theory entails only truths about what is rationally required, then one of those truths is that it is irrational to use that theory.

The resolute response: In the imprecise case, it amounts to a version of the permissivism about precise credences that I defended in Lecture 2.

Act-state dependence?

Worry: Aren't all our theories for rational choice in the presence of act-state dependence self-undermining?

A pragmatist argument against imprecise credences?

Recalling the definition of the pragmatic utility of a doxastic state from Lecture 1, we might take the pragmatic utility of a credal set to be the utility of the acts it leads you to perform when coupled with the correct decision theory. Then the fact about imprecise credences says that, for any genuinely imprecise credal set $\mathbf{P}$, there is a precise credence function P such that $\mathbf{P}$ prefers that P to $\mathbf{P}$.[13]

[13] Compare to impossibility results for scoring imprecise probabilities due to (Seidenfeld et al., 2012; Schoenfield, 2017; Mayo-Wilson & Wheeler, 2016).