

Arif Ahmed (University of Cambridge)

The point of rationality

What makes practical rationality a good idea? Hume's answer was that a rational person's means are suited to their ends. If Hume was right (and he was), then the transitivity of preference is not a requirement of rationality. Nor are Sen's principles alpha and beta. But gamma is.

Richard Bradley (London School of Economics and Political Science) & H. Orri Stefansson (Stockholm University)

Chance attitudes and Dynamic Consistency

Discussion of the relationship between dynamic consistency and the Sure-thing principle has figured prominently in recent debate over the rationality of the kind of ambiguity aversion displayed in the Ellsberg paradox; less so in the literature on the preference for fairness postulated by Diamond (1967). Yet both are instances of a preference for randomisation (respectively over events/states and people) that can lead to dynamically inconsistent sequences of choices and to an aversion to information. In this talk we argue that, for quite different reasons, these implications don't support that claim that the attitudes to chances exhibited in the Ellsberg and Diamond preferences are irrational. We then draw out the lessons of this for the much older debate over the significance of considerations of dynamic inconsistency for the rationality of the Allais preferences.

Till Grüne-Yanoff (Royal Institute of Technology, Stockholm)

A Building Code for Dynamic Preference Construction

Most people seem to have incomplete preferences. When the need for a preference comparison arises (when they must choose or express their opinion, for example) they often construct these preferences at that moment. Behavioral scientists have empirically identified such dynamic preference construction by showing that experimental subjects' preferences vary with contexts and are mutually inconsistent (Slovic 1995). In particular, preferences are shown to violate description and procedure invariance, well as independence of irrelevant alternatives. That preferences are dynamically constructed therefore are often considered part of the "cognitive bias" arsenal that many behavioral scientists refer to as evidence for limited human rationality – and as justification for paternalistic interventions.

Recently, however, this normative stance has come under scrutiny. Authors like McKenzie and Nelson (2003), Sher and McKenzie (2006, 2008), Geurts (2013), and Mandel (2014) have argued that dynamic preference construction is rational because it allows exploiting the informational value of contextual features. For example, listeners draw inferences from speaker's choice of frames, decisionmakers perceive defaults as implicit recommendations, or "irrelevant" alternatives serve as samples drawn from poorly known attribute distribution. Instead of marking (static) inconsistencies, constructed preferences should be seen as "dynamically coherent" in the view of these authors (McKenzie et al. 2018, 342). Unfortunately, though, the general notion of dynamic coherence is not further specified.

While I agree that it might sometimes be rational to delay preference construction in order to exploit such contextually conveyed information, I intuit that it is not always so. The task of my paper is to make this intuition precise, stating the conditions under which it is rational (i) to keep one's preferences incomplete, (ii) delay preference construction even though one is

committed to the description of a decision situation, (iii) adopt (heuristic) rules by which to construct preferences and (iv) construct or revise one's preferences in light of contextual features. The result will be a building code for dynamic preference construction that aims to explicate the idea of dynamic coherence.

Johan Gustafsson (University of Texas, Austin)

Causal Decision Theory and the Dutch-Book Argument for Conditioning

Causal decision theorists are vulnerable to a money pump if they update by conditioning when they learn what they have chosen. Nevertheless, causal decision theorists are immune to this money pump if they instead update by imaging on their choices and by conditioning on other things. I show that David Lewis's and Brian Skyrms's Dutch-book argument for conditioning does not work when you update on your choices. But a collective of causal decision theorists are still exploitable even if they start off with the same preferences and the same credences and will all see the same evidence. Evidential decision theorists who consistently update by conditioning are not exploitable in this way.

Peter Hammond (University of Warwick)

Rationality in Enlivened Decision Trees

In principle a decision-making agent's possible decisions and their uncertain consequences can be modelled and formally analysed using a decision tree. Yet except in relatively trivial cases, practicality puts bounds on the size and complexity of any tree that can be analysed properly. When the agent faces a sequence of decisions, the decision tree that the agent analyses may become "enlivened" between successive decisions by the addition of new nodes and new branches, possibly involving extensions to the set of states of the world, as well as enrichments of the consequence domain. My talk will first illustrate the concept of an enlivened decision tree by means of examples involving: (i) Homer's tale of Odysseus and the Sirens; (ii) José Luis Borges' "El Aleph", as a literary example of an unbounded model of the universe; (iii) births, marriages, and deaths; (iv) in economics, the related concepts of innovation and entrepreneurship in Joseph Schumpeter's 1911 "Theory of Economic Development"; (v) in decision theory, George Shackle's concept of "unexpected events", as well as Nassim Nicholas Taleb's of a "true" black swan, and the work of John Kay and Mervyn King on radical uncertainty; (vi) the games of Go and Chess. Nevertheless, one can model uncertainty about the ultimate retrospective value of any decision that the agent might make within whatever tree captures the current level of awareness. This allows the arguments of consequentialist decision theory that imply subjective expected utility maximization to be extended from standard to enlivened decision trees.

Conrad Heilmann (Erasmus University Rotterdam)

Inner conflict: shallow and deep

Cases of intertemporal weak will, such as procrastination, are associated with inner conflict. I advance a distinction between 'shallow' and 'deep' inner conflict that applies to such cases. Shallow inner conflict describes situations in which an individual promotes now over later. It is a time conflict. Deep inner conflict describes situations in which an individual has more than

one credible evaluative perspective which differ in recommendation. It is a value conflict. The distinction has a couple of payoffs. For one, the distinction can be used as a therapeutic tool. As such, it can help individuals with reflecting about what to do in cases that are commonly associated with intertemporal weak will. For another, distinguishing between shallow and deep inner conflict also has methodological advantages. The distinction effectively separates the time and value dimension in intertemporal weak will cases. This in turn, offers a new perspective with which to critically appraise philosophical and economic theories of intertemporal weak will.

Daniel Herrmann (University of California, Irvine)

Endogenizing Control

Decision theory describes how rational agents make decisions when they have control over things, using as input the agents' values and their beliefs about how the world works. But this leaves the objects of control to be entirely exogenously given. It would be very natural to desire an account of control that stems from the agent's own beliefs about how the world works. In this talk I will endogenize the notion of control. That is, I will show how we can extract, from an agent's attitudes, partitions over which it makes sense to say an agent views herself as having some degree of control. The core of my approach is a dynamic condition on an agent's degrees of belief I call desirability tracking. Desirability tracking will allow us to capture the standard situation in decision theory in which an agent has full control over a partition, as well as more general cases of partial control. After providing my desirability tracking account of control I will conclude by arguing that it reveals that act probabilities play a theoretical role in decision making.

Brian Hill (École des hautes études commerciales de Paris)

Dynamic consistency and ambiguity: A reappraisal

Doubtless the most important normative challenge to decision rules diverging from expected utility -- including so-called ambiguity rules -- concerns their implications in dynamic contexts. One common dynamic-choice-based argument against them purports to show that they are incompatible with the conjunction of two prima facie plausible principles: dynamic consistency and consequentialism. Dynamic consistency concerns preferences over contingent plans: so what counts are the contingencies the decision maker envisages -- and plans for -- rather than independently fixed contingencies, as implicitly assumed in standard formalisations. We show that an appropriate formulation of dynamic consistency to accommodate this point resolves the aforementioned conflict, hence undermining the criticisms of ambiguity models based on it. Moreover, it provides a reconceptualisation of dynamic choice under non-expected utility that neutralises many other dynamic- and sequential-choice-based arguments proposed in philosophy and economics. On the one hand, it provides a principled justification for the restriction to certain families of beliefs in applications of these models in dynamic choice problems. On the other hand, it supports a new analysis of the value of information under ambiguity, showing that decision makers may only turn down information if it has an opportunity cost, in terms of the compromising of information they had otherwise expected to receive.

Sebastian Krug (Leibniz Universität Hannover)

No escape from the cycle?

Classical Causal Decision Theory (CDT) has recently been criticised for generating costly conflicts between temporal stages of an agent in certain sequential choice problems. There have, accordingly, been some attempts to construct new broadly causalist decision theories which provide more diachronic consistency. The first part of my talk shows that the room for such improvement is rather constrained: Some more demanding consistency conditions can not be satisfied by any theory which qualifies as causal in a rather weak sense, as long as we do not adopt a strong (and implausible) version of resolute choice. Even several milder conditions – like not turning the agent into a money pump – can not be satisfied by such a theory if we uphold sophisticated choice. This partly builds on work by Dmitri Gallow, who has shown that causal decision theories generate something akin to cyclic preferences in certain examples. In the second part, I sketch my own attempt at improving the diachronic consistency of CDT without sliding to far into resolute choice. It combines (1) a new version of causal decision theory, with (2) a “mildly resolute” procedure for applying this theory to sequential problems. The former has similarities to proposals by Gustafson and by Rothfus, but differs in important regards. The resulting package satisfies several of the milder dynamic consistency conditions, which have been shown to be incompatible with causalism under sophisticated choice in the first part of the talk.

Richard Pettigrew (University of Bristol)

On Choosing how to Choose

It is our lot to face decisions when we have no certainty about which option from among those available to us will lead to the best outcome. How are we to make such choices? Expected utility theory is the most well-known and widely used account, but there are many alternatives available. Which should you use? Here's a natural approach to this problem: a decision theory is an account of rational means-ends reasoning, so it's natural to assess it by asking how well it performs in the role of getting you the ends that you in fact have. The only problem with this approach is that, in order to assess a decision theory or anything else as a means to your ends, we need an account of which means to your ends it is rational to use. And without a decision theory, we don't have that. Yet all is not lost, for this line of thinking nonetheless furnishes us with a test we can conduct on a theory of decision-making. We can ask of the theory: If I were to use you not only to make my normal day-to-day decisions, but also to make the higher-order decision about which decision theory to use, would you recommend yourself? If it would, we call it self-recommending. I'll argue it is a necessary but not sufficient condition on an adequate decision theory that it is self-recommending. I show that expected utility theory is self-recommending, but its most popular rivals are not.

Gerard Rothfus (The University of North Carolina at Chapel Hill)

Dynamic Consistency for Two Kinds of Moral Absolutism

Consider two logically independent moral doctrines that might justly be dubbed “absolutist”: *axiological absolutism* maintains that certain kinds of values (e.g., a person's life) categorically trivialize others (e.g., financial profit) for purposes of moral deliberation, while *motivational*

absolutism contends that certain kinds of values (again, perhaps, a person's life) can never be sacrificed as an instrumental means of securing ulterior ends, however valuable. Moral decision theorists have notably faulted axiological absolutism for apparently breeding potential dynamic inconsistencies in contexts of sequential choice under uncertainty but have left the dynamic consistency of motivational absolutism unexplored. This talk investigates the prospects for incorporating absolutist considerations of both sorts into a dynamically consistent moral decision theory, arguing that, while the desideratum of dynamic consistency does place constraints upon how we understand the content of axiological absolutism, both sorts of moral absolutism are ultimately compatible with the adoption of coherently implementable plans across time.

Toby Solomon (Ludwig-Maximilians-Universität, München)

Are my future choices DARC to me now?

An influential thesis has it that rational decision-makers cannot have credences over their present choices. Alan Hájek has memorably dubbed this the Deliberation Annihilates Reflexive Credences (DARC) thesis. An obvious question arises when we consider DARC in the context of dynamic decision-making: does the restriction on credences about my present choice extend to my future (or past) choices? To put it poetically: are my future choices DARC to me now? There are difficulties with both positive and negative answers. A positive answer will cause problems in cases where what I should do now is dependent on what I am most likely to do later—for example, when deciding whether to buy insurance for my own bad driving. While a negative answer leads to practical inconsistency in dynamic choice: my decisions, taken at different times, may work against each other even though I gain no new evidence and do not change my preferences. In this talk I will examine the prospects for developing a middle way for DARC in dynamic contexts that places some restrictions on my credences about future decisions, but does not rule them out entirely. Perhaps my future choices are only partially DARC to me now.

Wolfgang Spohn (University of Konstanz)

Reflexive Rationality and Sophisticated Choice

The first aim of the talk will be to briefly explain what hides behind the label “reflexive rationality”. At its base is a conception of strategies which makes strategies not depend on external states of affairs, but internally on envisaged changes of decision situations that may contain probabilities and utilities changed in arbitrary ways (and not only through information). The label thus stands for an apparently powerful program having many consequences, in particular also regarding game theory. However, these consequences remain in the abstract so far. Therefore, the second aim is to give substance to the program in at least in one respect, i.e., regarding dynamic choice. It will be argued that a crucial element in dynamic choice has been neglected so far, in particular in the prevailing decision rule of sophisticated choice: namely the agent’s evaluation of the envisaged changes of her decision situation as (more or less) favorable or unfavorable, an evaluation representable only within the reflexive perspective. This point will culminate in a more general decision rule, which I call reflective choice and which generalizes sophisticated choice.

Reuben Stern (Duke University, Durham)

The Chances of Choices

Meek and Glymour (1994) famously maintain that if we treat decision-theoretic options as interventions, then we can use evidential decision theory to vindicate causal dominance reasoning. Their basic idea is that this is guaranteed by a causal modeling axiom that implies that interventions are probabilistically independent of their non-effects -- namely, the Causal Markov Condition (CMC). But in previous work, I've raised two concerns for this argument. First, the CMC is not plausible when we interpret it as implying constraints on subjective probability functions, because there are cases involving causal uncertainty where it is rational for an agent to regard variables as causally independent but subjectively probabilistically dependent. Second, even absent causal uncertainty, the CMC doesn't imply that an agent should regard their intervention as probabilistically independent from its non-effects when the agent has "exotic evidence" -- i.e., evidence about some variable that they regard as causally downstream from their intervention. In this paper, I will argue that interventionists like Meek and Glymour can answer these challenges by adopting a conception of choice according to which there are significant constraints on the objective probabilities for decision-theoretic options (e.g., a conception of choice according to which the chance of every option must equal the chance of every other option). But I will also assess this line of response by considering whether it mandates obviously irrational choices in contexts where the agent's rational choice *prima facie* seems to depend on the agent's prediction(s) of their own choice(s).

Johanna Thoma (University of Bayreuth)

Some Mistakes are Irreducibly Diachronic

Behavioural welfare economics is concerned with agents who don't always display consistent, context-independent and stable preferences in their choice behaviour. There is controversy over whether perfect rationality always requires consistency, context-independence and stability. I argue that, whether or not we accept that it does, agents who violate these conditions sometimes make mistakes that are irreducibly diachronic: They perform sequences of choices that together unambiguously do not serve their subjective interests well, while no individual choice can be identified as a mistake. I then explore the implications of this thesis for behavioural welfare economics. For one, attempts at preference purification are often futile. Moreover, behavioural welfare economics should be more concerned with identifying mistakes of the diachronic type, as these may provide an additional basis for means paternalist intervention.

Kangyu Wang (London School of Economics and Political Science)

Sophisticated sophistication strategy in dynamic decision making

The sophistication strategy requires an agent to choose the option whose outcome is most desired by the agent when the choice is made while viewing their possible future choices as states of the world. It has two major problems. Firstly, it may seem counterintuitive when transformative experiences are involved. It sometimes seems intuitively rational to take future preferences into consideration when one knows that one will "become a different person". Secondly, sophistication may also seem counterintuitive because most people find that being resolute is rational, although resolution, on the face of it, seems "unsophisticated". The

sophisticated sophistication strategy gets all these intuitions right by redescribing the outcomes regarding the agent's desires, aversions and motivations. A desire, an aversion or a motivation, currently obtaining or to obtain in the future, can also be an object of another desire or aversion. A sophisticatedly sophisticated agent takes higher-order desires and aversions into account and in this way avoids those two problems.

Timothy Williamson (University of Oxford)

Retrospective Choice

I outline and defend a novel dynamic decision rule, *Retrospective Choice*. This rule is grounded in a simple idea: rational agents make choices in order to realise a future self who is happy to exist. In slightly more detail, this means that choices in a dynamic decision situation are evaluated based on the consequences they promote *by the lights of the final self who ends up with those consequences*. Once spelled out formally, Retrospective Choice shields agents with cyclic preferences from money pumps in deterministic decision situations.

I argue that Retrospective Choice rule is better (or at least more minimally) motivated than Resolute Choice. In particular, it is compatible with a parsimonious consequentialist framework that does away with the need for commitments, intentions, and similar features of diachronically extended agency. I also argue that it is extensionally superior to Sophisticated Choice. In particular, it shields agents with cyclic preferences from making inefficient choices in a well-known variant on the money pump due to John Cantwell. I conclude by sketching the implications of Retrospective Choice for decision-making when we face the possibility of preference change.