Justifying conditionalisation: conditionalisation maximizes expected epistemic utility

Hilary Greaves and David Wallace*

April 20, 2006

Abstract

According to Bayesian epistemology, the epistemically rational agent updates her beliefs by conditionalisation: that is, her posterior subjective probability after taking account of evidence X, p_{new}, is to be set equal to her prior conditional probability $p_{old}(\cdot|X)$. Bayesians can be challenged to provide a justification for their claim that conditionalisation is recommended by *rationality* — whence the normative force of the injunction to conditionalise?

There are several existing justifications for conditionalisation, but none directly addresses the idea that conditionalisation will be epistemically rational if and only if it can reasonably be expected to lead to epistemically good outcomes. We apply the approach of cognitive decision theory to provide a justification for conditionalisation using precisely that idea. We assign epistemic utility functions to epistemically rational agents; an agent's epistemic utility is to depend both upon the actual state of the world and on the agent's credence distribution over possible states. We prove that, under independently motivated conditions, conditionalisation is the unique updating rule that maximizes expected epistemic utility.

1 Introduction: Justifying conditionalisation

According to Bayesian orthodoxy, the ideal epistemic agent can be modelled as follows. The agent contemplates a set Ω of possible worlds. At every time, the agent's epistemic state can be represented by a probability function p over Ω (that is, probabilism holds). A learning event occurs when, for some subset X of Ω, the agent learns that the actual world is a member of X. On learning this, the agent updates her probability function by conditionalisation on X. That is,

Hilary Greaves, Department of Philosophy, Rutgers University, 26 Nichol Avenue, New Brunswick, NJ 08901-2882, USA. *Email:* hgreaves@rci.rutgers.edu.
David Wallace, Balliol College, Oxford OX1 3BJ, UK. *Email:* david.wallace@philosophy.ox.ac.uk.

her new credence function is related to her old by

$$p_{new}(\cdot) = p_{old}(\cdot|X),$$

where the conditional probability $p(A|B)$ is defined by

$$p(A|B) := \frac{p(A \wedge B)}{p(B)}.$$

Real epistemic agents are not (at least not quite) like this: Bayesian epistemology is a normative theory, rather than a purely descriptive one. But then Bayesians face the challenge: whence the normative force of the injunction to conditionalise?

One answer is 'it's *obvious* that conditionalisation is the epistemically rational way to update one's credence distribution'. This is true, but it is better if we can justify the obvious. Existing justifications for conditionalisation include: a Dutch Book argument (if you update your beliefs other than by conditionalisation, then a Dutch Book can be made against you [9]); an appeal to a Principle of Minimum Information (conditionalisation gives the posterior credence distributions that are 'closest' to your prior distribution while being consistent with the fact that you have just learnt X [13]); an appeal to a Principle of Reflection (Reflection entails conditionalisation [11]); and a symmetry argument ([10]:331-337). While these approaches have their interest and merits, none directly addresses the idea that conditionalisation will be epistemically rational if and only if it can reasonably be expected to lead to *epistemically good outcomes*.

This paper applies the approach of cognitive decision theory to provide a new justification for conditionalisation, based on precisely that idea. We assume that an epistemically rational agent always chooses that epistemic act that maximizes his expected epistemic utility, and we prove that, under independently motivated conditions, conditionalisation maximizes expected epistemic utility.

Section 2 is an introduction to the basic ideas of the cognitive decision theory we will use, including that of epistemic utility. We illustrate, by means of a toy example, how the fact that a given agent will always maximize expected epistemic utility is supposed to determine his choice of updating policy.

Section 3 contains the central claims of this paper. We note (section 3.1) that an agent faithfully represented by cognitive decision theory might (depending on the form of his utility function) be forbidden to hold some particular subset of possible credence distributions, on the grounds that the act of continuing to hold one of those credence distributions fails to maximize expected epistemic utility *calculated with respect to that same credence distribution*; that is, it may be that some credence distributions undermine themselves. We then prove (sections 3.2 and 3.3) that, of all possible belief-updating policies, conditionalisation maximizes expected epistemic utility provided only that the conditional probability distributions are not self-undermining in this sense.

There are two perspectives one might take on this result. First, *if* we regard it as a rationality constraint that the epistemic utility function must not forbid *a priori* any credence distribution in this way, then we will regard the proof as showing that for any epistemically rational agent, conditionalisation maximizes expected epistemic utility. Second, whether or not we accept that constraint as a necessary condition for epistemic rationality, it has been shown *for an arbitrary epistemic utility function* that the EU-maximizing agent conditionalises whenever his conditional posterior does not undermine itself. Since an agent whose conditional posteriors *do* undermine themselves obviously should *not* conditionalise, this is not only as strong an optimality proof as someone unwilling to accept constraints on utility functions can reasonably expect – it is also as strong an optimality proof as she could want.

Section 4 considers a few particular candidate epistemic utility functions, by way of illustration. Section 4.1 suggests a plausible epistemic utility function, and discusses how this particular utility function encodes a respect for epistemic values such as truth and verisimilitude. Section 4.2 discusses a *prima facie* intuitive, but on reflection less plausible, utility function that has been considered in the literature, according to which epistemic utility is *linearly related* to degree of belief in the truth. We discuss an objection to the approach of this paper: the objection that the possibility of utility functions such as this undermines the whole decision-theoretic approach to probabilist epistemic rationality.

Section 5 is the conclusion.

2 Cognitive decision theory

This section introduces the basic ideas of the cognitive decision theory we will use: states, probability distributions, epistemic acts, act availability, epistemic utility and expected epistemic utility. We explicate each of these notions below (section 2.1), and illustrate the theory throughout by means of a toy model of a simple cognitive decision process. Following this exposition, section 2.2 mentions, only to set aside, two closely related issues that we do *not* intend to address: the place (or lack of it) of cognitive decision theory in an 'all-things-considered' decision theory, and the relevance (or lack of it) of epistemic utility to the advisability of *gathering*, as opposed to epistemically responding to, evidence.

2.1 The framework of cognitive decision theory

Some cognitive states are, epistemically speaking, better than others. For example, it is (presumably) epistemically better to have higher credences in truths and lower credences in falsehoods. According to the cognitive decision-theoretic approach, epistemic rationality consists in taking steps that can reasonably be expected to bring about epistemically good outcomes.

Cognitive decision theory provides a framework in which the ideas of the preceding paragraph can be made precise and quantitative. The decision problems with which we will be concerned take the following form. The agent begins in some fixed belief state (that is, he holds some fixed initial credence distribution). He knows that he is about to receive some new piece of information, from among a fixed range of possibilities. Before receiving the information, he chooses an updating policy: that is, he specifies, for each of the possible pieces of new information, how he will change his credence distribution if that turns out to be the information that he does in fact receive. The decision he has to make is the choice of an updating policy.

> **EXAMPLE.** Mike has a coin. He is unsure as to whether or not it is a fair coin — specifically, he assigns 50% credence to its being fair — but he is (let us suppose) certain that *either* it is fair *or* it is weighted in such a way that the chances for outcomes (Heads, Tails) on a given toss are $(\frac{1}{4}, \frac{3}{4})$ respectively. The coin is about to be tossed; after observing the result of the toss, Mike will reassess his degrees of belief as to whether or not the coin is fair. He must decide in advance how the reassessment will proceed: which credence distribution he will move to if he sees heads, and which if he sees tails. We want to know how that decision should proceed.

The remainder of section 2.1 spells this out in more detail, in a framework of cognitive decision theory. (Cognitive decision theory is in many respects similar to ordinary, prudential decision theory; our framework is loosely based on that of Savage [8].)

States. The agent contemplates a set $\mathcal{S}$ of (mutually exclusive and jointly exhaustive) possible **states of the world**; he is unsure as to which element of $\mathcal{S}$ obtains. $\mathcal{S}$ can be thought of as a partition of the set of possible worlds.[1]

> For our toy example, the states might be as follows:
>
> $\mathcal{S} = \{s_{FH}, s_{FT}, s_{UH}, s_{UT}\}$, where
> s_{FH} : coin fair, outcome of toss is H
> s_{FT} : coin fair, outcome of toss is T
> s_{UH} : coin unfair, outcome of toss is H
> s_{UT} : coin unfair, outcome of toss is T.

Probability distributions. The agent does, however, have definite subjective degrees of belief as to which state obtains: his belief state is, at any time, represented by some probability distribution p over $\mathcal{S}$. We write $\mathcal{P}$ for the set of *all* probability distributions over $\mathcal{S}$.

The agent's prior. One particular probability distribution $p^* \in \mathcal{P}$ represents the agent's *prior* belief state — his belief state before learning the evidence, when he is making his cognitive decision.

> Mike's prior belief state is represented by the following probability distribution over $\mathcal{S}$:
>
> $$p^*(s_{FH}) = \tfrac{1}{4};\quad p^*(s_{FT}) = \tfrac{1}{4};\quad p^*(s_{UH}) = \tfrac{1}{8};\quad p^*(s_{UT}) = \tfrac{3}{8}.$$

[1] We will assume throughout that $\mathcal{S}$ is finite. This is merely for simplicity of exposition.

Experiments. An **experiment** is a situtaion in which the agent is to receive some new piece of information, from among a set of mutually exclusive and jointly exhaustive alternatives. Mathematically, an experiment is represented by a partition $\mathbf{E}$ of $\mathcal{S}$; we say that an experiment $\mathbf{E}$ is *performed* (for a particular agent) when the members of $\mathbf{E}$ become epistemically distinguishable from one another (for that agent).

Mike's experiment is given by $\mathbf{E} = \{H, T\}$, where $H = \{s_{FH}, s_{UH}\}$ and $T = \{s_{FT}, s_{UT}\}$.

Available acts. The set $\mathcal{A}_{\mathbf{E}}$ of **acts that are available given experiment E** is supposed to reflect the possible courses of (epistemic) action among which the agent must choose. An available epistemic act $a \in \mathcal{A}_{\mathbf{E}}$ is an assignment of a probability distribution to each piece of possible information $E_j \in \mathbf{E}$, with the intended interpretation that if $a(E_j) = p_j$, then p_j is the probability function that an agent performing act a would adopt as his credence distribution if he received the new information that the actual state was some member of E_j. (We make no assumption that the 'choice' between acts is a voluntary one.)

The acts that are available to Mike as a result of his experiment $\mathbf{E} = \{H, T\}$ are just those that assign one probability distribution over $\mathcal{S}$ to H, and another (not necessarily distinct) to T.

Acts. There is also a wider notion of epistemic act; the acts in this wider sense form a superset of the available acts $a \in \mathcal{A}_{\mathbf{E}}$. In the wider sense, an epistemic act $a' \in \mathcal{A}$ is an assignment of a probability function to every *state* $s \in \mathcal{S}$. The intended interpretation is that, if $a(s) = p_s$, then p_s is the probability function that an agent performing act a' would adopt as his credence distribution if state s in fact obtained.

An act $a' \in \mathcal{A}$ will fail to correspond to any *available* act $a \in \mathcal{A}_{\mathbf{E}}$ iff there is any pair of states $s_1, s_2 \in \mathcal{S}$ such that (i) s_1 and s_2 are members of the same element E_j of the partition $\mathbf{E}$, but also (ii) $a'(s_1) \neq a'(s_2)$. When this happens, a' is an act that the agent is not able to perform, since performing act a' would require the agent to respond to information that he does not have (hence our refusal to call such acts 'available').

Mike would be doing very well, epistemically speaking, if he were to perform the epistemic act $\tilde{a}$ corresponding to the instruction 'place credence unity in the true state', i.e.

$$\begin{aligned}\tilde{a}(s_{FH}) &= p_{FH}\\ \tilde{a}(s_{FT}) &= p_{FT}\\ \tilde{a}(s_{UH}) &= p_{UH}\\ \tilde{a}(s_{UT}) &= p_{UT},\end{aligned}$$

where, for each state s_i, $p_i(s_i) = 1$. However, this act is not available given only the experiment $\mathbf{E}$ defined above: Mike is going to receive certain information only about which side of the coin lands face up, so he cannot perform any act $a' \in \mathcal{A}$ with $a'(s_{FH}) \neq a'(s_{UH})$, or one with $a'(s_{FT}) \neq a'(s_{UT})$.

Given any available act $a \in \mathcal{A}_{\mathbf{E}}$, there is a natural way of identifying a with a particular act (in the wider sense), $a' \in \mathcal{A}$: simply set $a'(s) := a(E_j)$ whenever $s \in E_j$. Given this identification, we will sometimes slide between the two notations $a(s)$ and $a(E_j)$, for convenience.

Cognitive decision problem. A cognitive decision problem is specified by an experiment $\mathbf{E}$. To solve such a problem is to select an updating policy from $\mathcal{A}_{\mathbf{E}}$, the set of acts that are available given this experiment.

Mike's cognitive decision involves the decision of whether to commit himself to updating by conditionalisation from his prior p^* on the result of the coin toss, or by a particular rival updating policy **R** (given below) that has just popped into his head.[2]

Updating by conditionalisation from the prior p^* would lead to the following possible posteriors:[3]

$$\begin{array}{llllll} \textsf{\textbf{Cond}}(H) = p(\cdot|H) & =: p_H, & \text{where} & p_H(s_{FH}) & = & \frac{2}{3} \\ & & & p_H(s_{FT}) & = & 0 \\ & & & p_H(s_{UH}) & = & \frac{1}{3} \\ & & & p_H(s_{UT}) & = & 0; \\ \textsf{\textbf{Cond}}(T) = p(\cdot|T) & =: p_T, & \text{where} & p_T(s_{FH}) & = & 0 \\ & & & p_T(s_{FT}) & = & \frac{2}{5} \\ & & & p_T(s_{UH}) & = & 0 \\ & & & p_T(s_{UT}) & = & \frac{3}{5}. \end{array}$$

We stipulate that the alternative updating policy **R**, on the other hand, is as

[2]This is a simplification, of course. There are actually infinitely many acts that are available given $\mathbf{E}$; we restrict our attention, in this example, to just two available acts, for simplicity of exposition.

[3]Recall that we have defined an act (updating policy) as a function from states (or disjunctions of states) to probability distributions. According to this definition, conditionalisation from some prior p^* and conditionalisation from a distinct prior $q \neq p^*$ count as distinct updating policies, as do conditionalisation from prior p^* given an experiment $\mathbf{E}$ and conditionalisation from p^* given a distinct experiment $\mathbf{E}'$. Strictly, to pick out a unique act, we should therefore write $\mathbf{Cond}^{p^*}_{\mathbf{E}}$; we will shorten this to **Cond** since the prior p^* and experiment $\mathbf{E}$ will be held fixed throughout our discussion. This way of speaking has the advantage that the expected utility (with respect to some fixed probability distribution that may or may not equal the agent's prior) of an act will be independent of what the agent's prior was, and independent of which experiment is performed. It has the disadvantage that conditionalisation *simpliciter* does not count as a single updating policy, which is perhaps contrary to ordinary usage of the term 'updating policy'.

follows:

$$\begin{array}{llrcl} \mathbf{R}(H) = q_H, & \text{where} & q_H(s_{FH}) & = & \frac{1}{2} \\ & & q_H(s_{FT}) & = & 0 \\ & & q_H(s_{UH}) & = & \frac{1}{2} \\ & & q_H(s_{UT}) & = & 0; \\ \mathbf{R}(T) = q_T, & \text{where} & q_T(s_{FH}) & = & 0 \\ & & q_T(s_{FT}) & = & \frac{1}{4} \\ & & q_T(s_{UH}) & = & 0 \\ & & q_T(s_{UT}) & = & \frac{3}{4}. \end{array}$$

Cond and **R** are both *available* acts after the coin-flip, since each assigns the same probability distribution (**Cond**(H) and **R**(H) resp.) to s_{FH} as it does to s_{UH}, and the same probability distribution (**Cond**(T) and **R**(T) resp.) to s_{FT} as it does to s_{FT}. That is, neither **Cond** nor **R** makes the unreasonable requirement that Mike must take a different course of epistemic action depending (say) on whether the world, unbeknownst to Mike, happens to be in some fair-coin state or in some unfair-coin state; both updating rules require him to react only to information that the experiment will provide him with.

We offer no intuitive rationale for the rule **R**, and indeed we have none. The point is not that **R** has any intuitive plausibility whatsoever as a serious rival to conditionalisation, but rather that **R** is a course of epistemic action that an agent could in principle adopt. Our aim is to show that considerations of intuitive plausibility need not be invoked in order to outlaw **R**, because the inferiority of that updating policy will follow by calculation from the decision-theoretic model.

Epistemic utility functions. A given agent (we are assuming) holds a particular **epistemic utility function** — a function $U : \mathcal{S} \times \mathcal{P} \to \Re$ assigning a real number to each pair consisting of a state and a probability distribution. $U(s,p)$ represents the epistemic value ("epistemic utility") of holding credence function p when state s in fact obtains.[4]

Note that we allow our notion of utility to be *externalist* in the sense that we allow two pairs $< s,p >$, $< s',p >$, in which the agent is in the same cognitive state but a different state of the world in fact obtains, to be valued differently. This is to be expected since epistemic rationality may well value *truth*, over and above the subjective feelings associated with being in some given belief state.

Presumably, since he is a responsible epistemic agent, Mike attaches a high epistemic utility to having high degrees of beliefs in truths. In that case, his epistemic utility function might look something like this:

For arbitrary state $s \in \mathcal{S}$ and probability distribution p over $\mathcal{S}$,
$U(s,p) = -(1-p(s))^2 - \sum_{s' \neq s}(p(s'))^2$. (1)

[4] A different sort of cognitive decision theory (e.g. Levi [4], Maher [5]) focusses on cognitive acts that involve *accepting particular propositions*, rather than holding particular probability distributions. The domain of the epistemic utility function for such acceptance-based cognitive decision theories is the set of pairs $< s, A >$ of states and *propositions*: $< s, A >$ is to be read 'accepting proposition A when state s obtains'. We do not regard such theories as necessarily in competition with our own; they are engaged in a different project.

Heuristically, we can see that this utility function 'values truth' because it equates the utility of holding credence function p when state s obtains to the sum of two terms, the first of which increases with increasing degree of belief in the true state s, and the second of which increases with decreasing degrees of belief in false states s'. (Note the minus signs.)

Expected epistemic utility. We assume[5] that the rational epistemic agent always performs that act that has the highest expected epistemic utility (with respect to the agent's prior p^*) of all available acts, where the expected epistemic utility of an act a (with repsect to probability distribution p) is given by

$$EU^p(a) = \sum_{s \in \mathcal{S}} p(s) \cdot U(s, a(s)) \tag{2}$$

Using the alternative notation mentioned above, in which acts are defined on experiments rather than directly on $\mathcal{S}$, we can also write the expected utility of a as

$$EU^p(a) = \sum_{E_j \in \mathbf{E}} \sum_{s \in E_j} p(s) \cdot U(s, a(E_j)), \tag{3}$$

where $\mathbf{E}$ is the experiment on which the act a is defined.

Iff an act a maximizes expected epistemic utility given an experiment $\mathbf{E}$ (that is, if it has at least as high an expected epistemic utility as any other act in $\mathcal{A}_{\mathbf{E}}$), we say that a is **optimal** (given $\mathbf{E}$). Iff a is the *unique* optimal act, we say that a is **strongly optimal** (given $\mathbf{E}$). (We will often leave 'given $\mathbf{E}$' implicit.)

Being an epistemically rational agent, Mike will choose whichever updating policy has the higher expected epistemic utility. To see which policy this is, we evaluate the EU of each policy using Mike's prior p^* and his epistemic utility function U, as follows:

Expected epistemic utility of updating by conditionalisation from prior p^* given experiment **E**,

$$\begin{aligned}
EU^{p^*}(\mathbf{Cond}) &= \textstyle\sum_{s \in \mathcal{S}} p^*(s) \cdot U(s, \mathbf{Cond}(s)) \\
&= p^*(s_{FH}) \cdot U(s_{FH}, p_H) + p^*(s_{FT}) \cdot U(s_{FT}, p_T) \\
&\quad + p^*(s_{UH}) \cdot U(s_{UH}, p_H) + p^*(s_{UT}) \cdot U(s_{UT}, p_T) \\
&= \tfrac{1}{4}(-(1-\tfrac{2}{3})^2 - (\tfrac{1}{3})^2) + \tfrac{1}{4}(-(1-\tfrac{2}{5})^2 - (\tfrac{3}{5})^2) \\
&\quad + \tfrac{1}{8}(-(1-\tfrac{1}{3})^2 - (\tfrac{2}{3})^2) + \tfrac{3}{8}(-(1-\tfrac{3}{5})^2 - (\tfrac{2}{5})^2) \\
&= -0.0479.
\end{aligned}$$

[5]Without providing a justification. Perhaps the justification is that a representation theorem shows that any agent whose preferences over acts satisfy certain axioms must be representable by some utility function U and some probability distribution p such that he always prefers an act a to another b iff $EU^p(a) > EU^p(b)$, with EU^p as defined in (3). Or perhaps the justification is simply that the injunction to maximize EU seems intuitively plausible and gives intuitively plausible results. We won't go into this issue; we take as a premise, for the purposes of this paper, that there is *some* adequate justification.

On the other hand, the expected epistemic utility of adopting the alternative policy **R** is given by

$$\begin{array}{rcl} EU^{p^*}(\mathbf{R}) & = & \sum_{s\in\mathcal{S}} p^*(s)\cdot U(s,\mathbf{R}(s)) \\ & = & p^*(s_{FH})\cdot U(s_{FH},q_H)+p^*(s_{FT})\cdot U(s_{FT},q_T) \\ & & +p^*(s_{UH})\cdot U(s_{UH},q_H)+p^*(s_{UT})\cdot U(s_{UT},q_U) \\ & = & \frac{1}{4}(-(1-\frac{1}{2})^2-(\frac{1}{2})^2)+\frac{1}{4}(-(1-\frac{1}{4})^2 \\ & & -(\frac{3}{4})^2)+\frac{1}{8}(-(1-\frac{1}{2})^2-(\frac{1}{2})^2)+\frac{3}{8}(-(1-\frac{3}{4})^2-(\frac{1}{4})^2) \\ & = & -0.120. \end{array}$$

Since $EU^{p^*}(\mathbf{Cond}) > EU^{p^*}(\mathbf{R})$, Mike will choose to update by **Cond** rather than by the alternative policy **R**.

2.2 Disclaimers

Before proceeding further, we flag two issues that are closely related to the spirit of the decision-theoretic approach, but that we do not (and need not) tackle in this paper: the place (or lack of it) of cognitive decision theory in an 'all-things-considered' decision theory, and the relevance (or lack of it) of epistemic utility to the advisability of *gathering*, as opposed to epistemically responding to, evidence.

Epistemic vs. all-things-considered utility. There is a sense in which a particular belief state can have a high 'utility' without it being at all *epistemically* rational to pursue that state. A high degree of belief in the existence of a god, for instance, may make me happy and make my life go better in all sorts of ways — and thus be of high 'utility' in a prudential sense — but yet, if there is no god then such a belief state is, *epistemically* speaking, undesirable, and if the evidence for its truth is slim then it is *epistemically* irrational to hold such beliefs. We fully accept that there is also this prudential sense of 'utility', and that the demand to maximize *prudential* utility may conflict with the demands of epistemic rationality. But this does not entail that considerations of *epistemic* utility cannot hope to account for epistemic rationality. An *epistemic* utility function, such as those we work with in this paper, is concerned only with the *epistemic* desirability of belief states. Where epistemic desirability diverges from 'prudential' or 'all-things-considered' desirability, an epistemic utility function tracks only the former. Whether and how an epistemic utility function plays a role in any integrated theory of the 'all-things-considered' rationality of epistemic acts is an important question, but not one we are concerned with here.

Epistemic vs. non-epistemic acts. It has been argued that a decision-theoretic framework — cognitive ([6], [2]:127-9) or otherwise ([1]) — can be invoked to *justify experimentation*. However that may be, that is not the project we are engaged in here. Our application of cognitive decision theory is concerned with *purely epistemic acts*; the act of performing a particular experiment, however epistemically motivated, is a non-epistemic act. That is, we assume only that, *given* that one has received a given piece of evidence (against one's will or otherwise), epistemic rationality requires that one then perform the epistemic act of altering one's belief state in the manner that, in the light of that evidence, maximizes expected epistemic utility. Whether and how a cognitive decision theory for epistemic acts could be integrated into a satisfactory decision theory, paying due respect to epistemic goods, for choices among non-epistemic acts (including evidence-gathering acts) is an important question, but, again, not one that we are concerned with here.

3 Conditionalisation and maximization of expected epistemic utility

This section contains our claims in defense of conditionalisation. We proceed in three steps. In section 3.1 we define the notion of a *constant act*, and a relation we call *recommendation* between probability distributions. We note an important consequence of our assumption that an epistemically rational agent always chooses the epistemic act that maximizes expected epistemic utility: for some utility functions, there exist probability distributions that the ideal agent is forbidden to hold on the grounds that they fail to 'recommend' themselves. In section 3.2 we use the notion of recommendation to define a class $\mathcal{QC}$ of epistemic acts, the *quasi-conditionalising* acts, and we prove (for an arbitrary epistemic utility function) that each act in $\mathcal{QC}$ maximizes expected epistemic utility. In section 3.3 we characterize (Corollary 2) a set of epistemic utility functions for which *conditionalisation* is optimal. We also prove (Corollary 1; again for an arbitrary epistemic utility function) that in any case conditionalisation is optimal if it is even *coherent*, in the sense that the probabilities conditionalisation would have the agent adopt are not ones that his own utility function forbids him ever to hold.

3.1 Constant acts, recommendation, self-recommendation and stable utility functions

Constant acts. We will have particular interest in the **constant acts**: those acts that instruct the agent to adopt the *same* probability distribution as his credence function, regardless of which state obtains.

The expression for the expected utility of a constant act takes a particularly simple form. For arbitrary $q \in \mathcal{P}$, let k_q denote the constant act that assigns q to all states, for arbitrary $q \in \mathcal{P}$. The expected epistemic utility of a constant act k_q, calculated with respect to the probability function p, is given by

$$EU^p(k_q) = \sum_{s \in S} p(s) \cdot U(s, q). \quad (4)$$

The recommendation relation between probability functions. The notion of the epistemic utility of a constant act raises an interesting issue that is important to our present project. All constant acts are, of course, always *available* in the sense explicated in section 2.1: one does not need to be receiving any new information about what the world is like in order to perform the act 'jump to credence distribution q, regardless of what the world is like'. (To put this point another way: the trivial experiment, $\{\mathcal{S}\}$, is always being performed, and all constant acts are members of $\mathcal{A}_{\{\mathcal{S}\}}$). So, at every time, the agent is to regard all constant acts as available options between which he can choose.

This ever-present availability of all constant acts has an interesting consequence within our cognitive decision theory. Recall (from section 2.1) that we are assuming that the epistemically rational agent always performs that available act that maximizes expected epistemic utility. Therefore, an ideally rational agent is able to hold a probability distribution p as his credence distribution only if the corresponding constant act k_p, *by the lights of p itself*, maximizes expected epistemic utility — that is, only if $(\forall q \in \mathcal{P})(EU^p(k_p) \geq EU^p(k_q))$. If this condition fails, the minute the agent found himself holding p, he would be compelled to move to some other distribution q that maximized expected epistemic utility calculated with respect to p — which is to say that an ideally rational agent could not hold p in the first place, even for a moment.

We make the following definitions:

- Say that p **recommends** q (write $p \xrightarrow{R} q$) iff, when the only available acts are the constant acts, k_q maximizes expected utility calculated with respect to p — that is, iff $\forall r \in \mathcal{P}, EU^p(k_q) \geq EU^p(k_r)$.
- Iff, in addition, p recommends no distribution distinct from q, say that p **strongly recommends** q.
- Iff p recommends p, say that p **is self-recommending**.
- Iff, in addition, p recommends no distribution distinct from p, say that p is **strongly self-recommending**.
- Iff p is not self-recommending, say that p is **self-undermining**.

An epistemic utility function therefore induces a rich structure on the set $\mathcal{P}$ of possible credence functions. The structure of the recommendation relations between these possible credence functions, for a fixed epistemic utility function, can be represented by a directed graph. This is illustrated, for a small subset of the set $\mathcal{P}$ of probability functions on $\mathcal{S}$, in Figure 1.

Domain of stability of an epistemic utility function. Clearly, the extension of the recommendation relation depends on the utility function U. We can thus classify utility functions based on the structure of the recommendation relations they induce:

- Say that U is **everywhere stable** iff, according to U, every probability distribution is self-recommending.

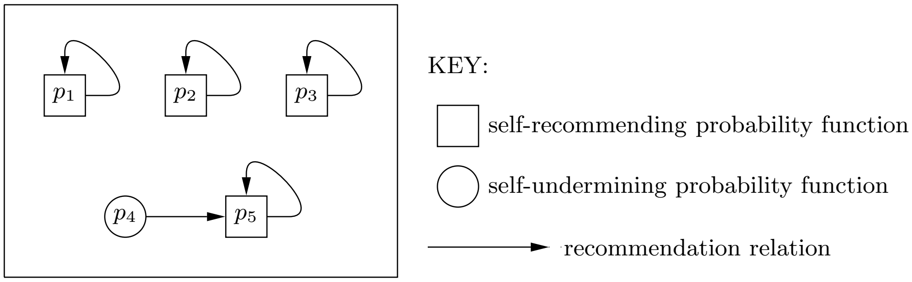

Figure 1: Graphical representation of a possible structure of recommendation relations between five probability functions $p_1, \ldots, p_5 \in \mathcal{P}$. In this example, the epistemically rational agent would not hold p_4 as his credence function for any finite length of time, since the expected epistemic utility, calculated with respect to p_4, of jumping to p_5 (performing the constant act k_{p_5}) exceeds that of remaining at p_4 (performing the constant act k_{p_4}) — since, in other words, p_4 fails to recommend itself but does recommend p_5.

- Say that U is **everywhere strongly stable** iff, according to U, every probability distribution is strongly self-recommending.
- Say that U is **somewhere stable** iff, according to U, some probability distributions are self-recommending and others are self-undermining.
- Say that U is **nowhere stable** iff, according to U, every probability distribution is self-undermining.

These four possibilities are illustrated in Figure 2.

We now consider the following question: which of these types of utility function might an epistemically rational agent hold?

Consider, first, Figure 2(d). This represents the recommendation structure induced by a nowhere stable utility function: no credence distribution recommends itself. In other words, whatever credence function the agent holds, he should at that same time consider some *other* credence function to be epistemically better *by the lights of his current credence function*. An agent who maximized expected epistemic utility with respect to a nowhere stable utility function would thus suffer from an epistemic version of the 'grass is always greener on the other side of the fence' syndrome. As a result, he would not be able to hold any given credence function for any finite interval of time, even in the absence of new information. This is pathological.

Next, consider figure 2(c). This represents a *somewhere stable* utility function: p_9, p_{10} and p_{12} are self-recommending, but p_{11} is self-undermining. At first sight, such somewhere stable utility functions (perhaps) also seem to be pathological — the notion of an ideal agent who held a somewhere stable utility

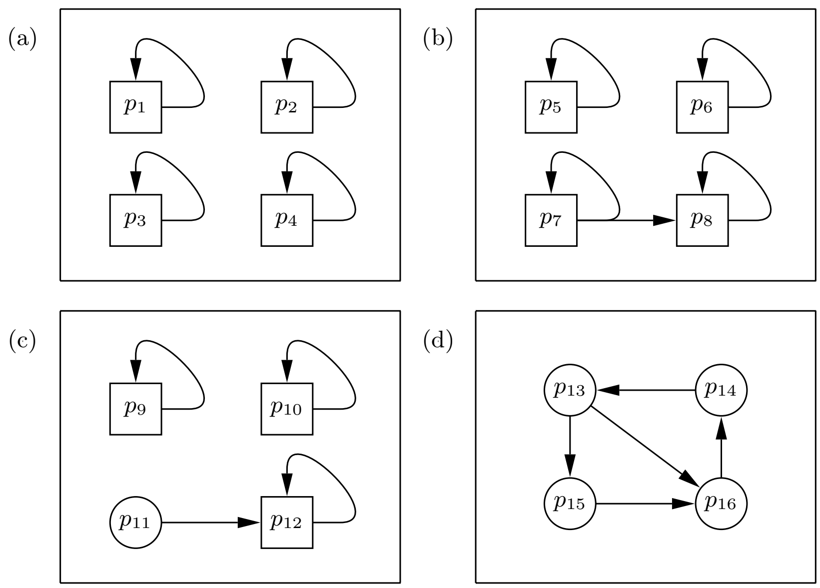

Figure 2: Recommendation structures induced by various types of epistemic utility functions: (a) everywhere strongly stable; (b) everywhere stable but not everywhere strongly stable; (c) somewhere stable; (d) nowhere stable.

function but also held (for any finite period of time) a credence distribution that that utility function deemed self-undermining would be similarly contradictory. A moment's reflection, however, shows that the problem here, if there is one, falls short of outright contradiction. Credence functions that fail to recommend themselves are 'forbidden points' in the space $\mathcal{P}$ of all possible credence functions. (The sense in which they are 'forbidden' is that it is not consistent with maximization of expected epistemic utility to remain at such a point for a finite length of time.) Nowhere stable utility functions, we saw, were pathological because they rendered *all* credence functions thus forbidden. We do not have this problem, though, in the case of *somewhere* stable utility functions. Provided that there are still *some* probability functions that are not 'forbidden' in this sense — provided, that is, that the utility function is not *nowhere* stable — there are ways for the agent to exhibit perfectly normal-looking epistemic behavior, consistent with maximization of expected epistemic utility.

Turn now to figures 2(a) and 2(b). The distinction between (merely) everywhere stable utility functions on the one hand, and those that are everywhere *strongly* stable on the other, also deserves comment. An agent who holds a utility function that is everywhere stable but fails to be everywhere *strongly* stable (as in figure 2(b)) may find himself with a choice between constant acts that are equally good by his own lights: if he currently holds credence distribution p_7, there is a distinct credence distribution p_8 such that, by the lights of p_7 itself, k_{p_7} and k_{p_8} are of equal (and optimal) expected epistemic utility. When this occurs, the agent *can* stick to his current distribution p_7, but it will be equally consistent with ideal rationality if he chooses to move to p_8 on a whim. An agent whose utility function is everywhere *strongly* stable (figure 2(a)), on the other hand, never has any such free options; he must always, unless new evidence comes along, stick to his current credence distribution. (This distinction will be of some importance in section 3.3.)

3.2 Theorem: Quasiconditionalisation maximizes expected epistemic utility

From this point onwards, we will assume that the agent's prior probabilities $\{p^*(E_j) : E_j \in \mathbf{E}\}$ are all non-zero.[6]

We give names to two updating policies in which we will have particular interest:

- *Conditionalisation from prior p^* given experiment $\boldsymbol{E}$* ($\mathbf{Cond}_{\mathbf{E}}^{p^*}$, or **Cond** for short) is defined (as usual) by

  Cond: For all $E_j \in \mathbf{E}, \mathbf{Cond}(E_j) = p^*(\cdot|E_j)$.

[6] This is for simplicity of exposition. Without this assumption, instead of working directly with available *acts* in the definitions and proofs that follow, we would work with equivalence classes of available acts, under the equivalence relation given by

$$R_1 \sim R_2 \text{ iff, for all } E_j \text{ such that } p^*(E_j) > 0,\ R_1(E_j) = R_2(E_j).$$

(Clearly, any two acts that bear this equivalence relation to one another have the same expected epistemic utility with respect to p^*.)

- *Quasi-conditionalisation from prior* p^* *given experiment* $\boldsymbol{E}$ ($\mathcal{QC}^{p^*}_{\mathbf{E}}$, or $\mathcal{QC}$ for short) is, in general, a set of updating policies rather than a single updating policy. It is defined as follows:

  Quasi-conditionalisation: For any available act $\mathbf{Q} : \mathbf{E} \to \mathcal{P}$, say that $\mathbf{Q}$ is a **quasi-conditionalising act** ($\mathbf{Q} \in \mathcal{QC}$) iff, for all $E_j \in \mathbf{E}$, $p^*(\cdot|E_j) \xrightarrow{R} \mathcal{QC}(E_j)$.

  That is, the quasiconditionalising acts are those acts according to which the agent, on receiving evidence E_j, moves, not necessarily to the conditional probability $p^*(\cdot|E_j)$, but to some probability distribution q that *is recommended by* that conditional probability (where 'recommendation' is as defined in section 3.1 above). (The reason that there may in general be more than one quasiconditionalising act is that $p^*(\cdot|E_j)$ may in general simultaneously recommend more than one distribution.) Write $\mathcal{QC} \subset \mathcal{A}_{\mathbf{E}}$ for the set of all quasi-conditionalising acts.

First, we prove that the optimal acts are exactly the quasi-conditionalising acts.[7]

Theorem. *Of all acts that are available given an experiment* $\boldsymbol{E}$*, all and only quasi-conditionalising acts are optimal. That is,*

$$\forall \mathbf{Q} \in \mathcal{QC}, \forall \mathbf{R} \in \mathcal{A}_{\mathbf{E}}, EU^{p^*}(\mathbf{Q}) \geq EU^{p^*}(\mathbf{R}),$$

with equality iff $\mathbf{R}$ is also a quasi-conditionalising act.

Proof. The expected utility of adopting an arbitrary updating policy $\mathbf{R} \in \mathcal{A}_E$ is given by

$$\begin{aligned}
EU^{p^*}(\mathbf{R}) &\equiv \sum_{s \in S} p^*(s) \cdot U(s, \mathbf{R}(s)) && (5)\\
&\equiv \sum_{E_j \in \mathbf{E}} \sum_{s \in E_j} p^*(s) \cdot U(s, \mathbf{R}(E_j)) && (6)\\
&\equiv \sum_{E_j \in \mathbf{E}} \sum_{s \in E_j} p^*(s \wedge E_j) \cdot U(s, \mathbf{R}(E_j)) && (7)\\
&\equiv \sum_{E_j \in \mathbf{E}} p^*(E_j) \cdot \left(\sum_{s \in E_j} p^*(s/E_j) \cdot U(s, \mathbf{R}(E_j)) \right) && (8)\\
&\equiv \sum_{E_j \in \mathbf{E}} p^*(E_j) \cdot \left(\sum_{s \in \mathcal{S}} p^*(s/E_j) \cdot U(s, \mathbf{R}(E_j)) \right) && (9)\\
&\equiv \sum_{E_j \in \mathbf{E}} p^*(E_j) \cdot EU^{p^*(\cdot|E_j)}(k_{\mathbf{R}(E_j)}). && (10)
\end{aligned}$$

[7] A form of the proof very similar to the one we present here was independently discovered by Frank Arntzenius, at the same time as us.

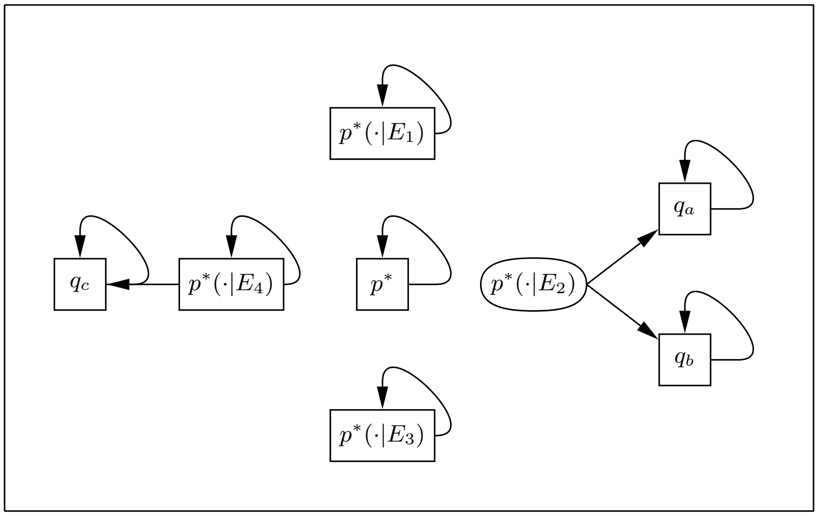

Figure 3: Diagram to illustrate the concept of a quasi-conditionalising act. For an available act $\mathbf{Q}$ to count as a quasi-conditionalising act, it has to be the case that, for all possible outcomes E_j of the experiment $\mathbf{E}$, the conditional probability $p^*(\cdot|E_j)$ recommends the credence function given by $\mathbf{Q}(E_j)$. The diagram represents a situation in which the partition $\mathbf{E}$ has four elements $\{E_1, E_2, E_3, E_4\}$; two of the conditional probabilities $(p^*(\cdot|E_1), p^*(\cdot|E_3))$ are strongly self-recommending, while the other two $(p^*(\cdot|E_2), p^*(\cdot|E_4))$ each recommend two probability functions. In this particular case, there are four quasi-conditionalising acts. For example, the following act $\mathbf{Q}'$ is a quasi-conditionalising act: $\mathbf{Q}'(E_1) = p^*(\cdot|E_1), \mathbf{Q}'(E_2) = q_b, \mathbf{Q}'(E_3) = p^*(\cdot|E_3), \mathbf{Q}'(E_4) = p^*(\cdot|E_4)$. Conditionalisation itself is *not* a quasi-conditionalising act in the situation depicted, since, here, $p^*(\cdot|E_2)$ is self-undermining.

Substituting a quasi-conditionalising act $\mathbf{Q} \in \mathcal{QC} \subset \mathcal{A}_\mathbf{E}$ for $\mathbf{R}$ in line (10) yields an expression for the expected utility of adopting that act $\mathbf{Q}$ as one's updating rule:

$$EU^{p^*}(\mathbf{Q}) = \sum_{E_j \in \mathbf{E}} p^*(E_j) \cdot EU^{p^*(\cdot|E_j)}(k_{\mathbf{Q}(E_j)}). \tag{11}$$

But, by the definition of quasi-conditionalisation, we have, for all available acts $\mathbf{R} \in \mathcal{A}_\mathbf{E}$, all quasi-conditionalising rules $\mathbf{Q} \in \mathcal{QC}$ and all events $E_j \in \mathbf{E}$,

$$EU^{p^*(\cdot|E_j)}(k_{\mathbf{Q}(E_j)}) \geq EU^{p^*(\cdot|E_j)}(k_{\mathbf{R}(E_j)}). \tag{12}$$

Combining (10), (11) and (12) and noting that the coefficients $p^*(E_j)$ are all nonnegative, we have

$$\forall \mathbf{Q} \in \mathcal{QC}, \forall \mathbf{R} \in \mathcal{A}_E, EU^{p^*}(\mathbf{Q}) \geq EU^{p^*}(\mathbf{R}). \tag{13}$$

Finally, we prove that this inequality is strict unless $\mathbf{R}$ is also a quasi-conditionalising act, as follows.

If $\mathbf{R}$ is not a quasi-conditionalising act, then there is some $E_k \in \mathbf{E}$ such that $EU^{p^*(\cdot|E_k)}(k_{\mathbf{Q}(E_k)}) > EU^{p^*(\cdot|E_k)}(k_{\mathbf{R}(E_k)})$. Since the terms in the summation (10) for $EU^{p^*}(\mathbf{R})$ and those in the summation (11) for $EU^{p^*}(\mathcal{QC})$ can then be paired off in such a way that for each pair, the term in $EU^{p^*}(\mathbf{R})$ is no greater than that in $EU^{p^*}(\mathcal{QC})$, and there is at least one pair such that the term in $EU^{p^*}(\mathbf{R})$ is strictly less than that in $EU^{p^*}(\mathcal{QC})$, it follows that $EU^{p^*}(\mathbf{Q}) > EU^{p^*}(\mathbf{R})$.

□

It is striking that this theorem involves *no assumptions whatsoever* about the nature of the utility function — yet we seem (at first sight) to have given a name to a particular set of epistemic acts (viz. $\mathcal{QC}$) and proved that every act in that set is optimal.

If this were really what we had done, it should arouse puzzlement in anyone acquainted with 'ordinary' (i.e. prudential, non-cognitive) decision theory — it is a familiar point from that theory that tastes are encoded in the utility function, so that one cannot prove anything about the EU of a given act *without* constraining the utility function. (There is no hope, for instance, of proving *from decision theory alone* that a rational agent tries to avoid being eaten by alligators; I may be perfectly 'rational', in the sense that I satisfy the axioms of decision theory, but happen to *like* being eaten by alligators, and accordingly assign high utility to situations in which I receive such treatment.) But our above 'first-sight' gloss on the content of our theorem is, of course, not quite correct. We have *not* shown, in the absence of any information about the utility function, that *some particular act* is optimal. This is because, in the absence of information about the utility function, we have no idea what the recommended probabilities $\{\mathbf{Q}(E_j)\}_{E_j \in \mathbf{E}}$ are. In other words, while we know (without knowing anything about the utility function) that all acts that meet the definition of $\mathcal{QC}$

are optimal, we do *not* know which acts (i.e., which functions from $\mathbf{E}$ to $\mathcal{P}$) those are.

3.3 Corollaries: When conditionalisation is optimal

The circumstances under which *conditionalisation* (as opposed to 'mere' quasi-conditionalisation) is optimal are brought out by the following two corollaries to our theorem.

Corollary 1. *Conditionalisation is optimal for a given experiment* $\boldsymbol{E}$ *iff the conditional probabilities* $\{p^*(\cdot|E_j) : E_j \in \boldsymbol{E}\}$ *are all self-recommending. Iff, further, these conditional probabilities are all strongly self-recommending, then conditionalisation is strongly optimal.*

Proof. Iff the conditional probabilities $\{p^*(\cdot|E_j) : E_j \in \mathbf{E}\}$ are all self-recommending, then conditionalisation is a quasi-conditionalising act. Iff, in addition, these conditional probabilities are all strongly self-recommending, then conditionalisation is the only quasi-conditionalising act. Corollary 1 is therefore immediate from the above theorem.

□

Corollary 1 establishes that conditionalisation is optimal whenever the conditional probabilities are self-recommending. Now, one who hoped to justify conditionalisation within a decision-theoretic framework really could not want stronger support from the mathematics, for the following reason. If the conditional probabilities are *not* self-recommending, conditionalisation is obviously not even a live option for our agent — for then, conditionalisation advises him to move to probabilities that he never would be able to hold, 'no matter how he arrived at them', while remaining an expected utility maximizer. (Compare our discussion of self-recommendation in section 3.1.) It would be somewhat worrying if our proof insisted *even then* that conditionalisation was optimal. So, corollary 1 establishes that conditionalisation is optimal whenever conditionalisation is even a live option. We stress that this follows from the decision theory alone, with no constraints on the form of the epistemic utility function. This is our first result in support of the normative status of conditionalisation.

Our second corollary concerns a second gloss we might put on our result, if we are prepared to accept normative constraints on the form of the epistemic utility function:

Corollary 2. *If the agent's epistemic utility function U is everywhere stable, then conditionalisation is optimal. If U is everywhere* strongly *stable, then conditionalisation is strongly optimal.*

Proof. This is an immediate consequence of Corollary 1. □

If it is a *rationality constraint* that one's epistemic utility function be everywhere strongly stable (so that one's utility function alone does not preclude holding any particular probability distribution, and always advises one strictly to stick to one's current credence distribution until and unless new evidence comes along), then Corollary 2 demonstrates that, for any rational agent, conditionalisation is the unique updating policy that maximizes expected epistemic utility. This would be a second statement in favor of conditionalisation. We find this rationality constraint plausible, but we offer no argument for it here.[8] (The reader may or may not find that she accepts it without argument.) If the constraint is not accepted, our categorical claims are restricted to those we drew above from Corollary 1.

4 A plausible epistemic utility function

So far, we have for the most part focussed on certain abstract features (everywhere/somewhere and weak/strong stability) of the epistemic utility function; such abstract features have sufficed to state and prove our claims. However, in order better to understand what is going on, we need to consider what a plausible utility function exhibiting some of these features might actually look like. In section 4.1 we take a brief look at one class of plausible (everywhere strongly stable) epistemic utility functions, and consider how functions in that class could encode various epistemic values. In section 4.2 we comment briefly on a particular somewhere stable utility function, the 'linear utility function', that has appeared in the literature. We answer an objection that the possibility of somewhere stable utility functions undermines the whole decision-theoretic approach.

4.1 An everywhere strongly stable utility function

Consider the following utility function schema:

General quadratic utility function : $U_{GQ}(s,p) = -\sum_{X\subseteq\mathcal{S}} \lambda_X(\chi_X(s)-p(X))^2$,

where χ_X is the characteristic function of the set X (that is, $\chi_X(s)$ is 1 if $s \in X$ and zero otherwise), and the λ_X are constant coefficients. (This is a generalization of the utility function (1) we used in our toy model in section 2.1.) We will now briefly discuss how epistemic utility functions of this form do justice to various epistemic norms.

[8]Wayne Myrvold has pointed out (personal correspondence) that the fact that we do have several other arguments to the effect that (within the domain of applicability of Bayesian modelling) conditionalisation is always rationally required is relevant here. (Some of these arguments were cited in section 1.) The expected-utility approach with no rule against somewhere stable utility functions, since it permits updating rules other than conditionalisation, is in tension with those results. If any of those other arguments is sound (a question we have not addressed), it may also contain the seeds of an explanation *from the perspective of cognitive decision theory* of how and why somewhere stable utility functions should be disallowed. We have not pursued this line of research.

Stability. For arbitrary choices of the coefficients λ_X, U_{GQ} is everywhere strongly stable. [9] [10]

A concern for truth. In the first instance, U_{GQ} favors placing credence 1 in the true state. *This* is a property of any stable utility function[11], and *a fortiori* of U_{GQ}. More generally (and less rigorously), other things being equal, U_{GQ} favors increasing one's credence in the true state. (We can see the latter by noting that $U_{GQ}(s, p)$ is always an increasing function of $p(s)$, and always a decreasing function of $p(s')$ for $s' \neq s$.)

Discriminating among falsehoods: taking account of verisimilitude. A different stable epistemic utility function (viz. $U(s, p) = \log p(s)$) encodes a sort of *epistemic perfectionism*: according to that utility function, epistemic utility depends *only* on credence in the true state. Such perfectionism may, on occasion, be appropriate. But, often, we will want instead to judge one credence distribution as epistemically better than another even when both assign the same degree of belief to the true state, on the grounds that the first concentrates its remaining credence among (false) states that are *closer to the truth* than does the second. Our sample schema U_{GQ} can take account of the value of verisimilitude, by a judicious choice of the coefficients λ_X: we simply assign high λ_X when X is a set of 'close' states.

Informativeness. Discussions of the epistemic desirability of holding *informative* or *contentful* beliefs are important in *acceptance*-based epistemologies, as opposed to the purely probabilist epistemology under consideration here – given that you're going to accept (say) some true proposition, it is epistemically better to adopt a more informative one, i.e. a stronger one. In the probabilist case, however, the epistemic value of informativeness is already captured by attaching epistemic value to truth and to verisimilitude – an agent will do better in terms of truth-credence and verisimilitude by peaking his probability

[9] Proof: in each case, use Lagrange multipliers to extremize the expected utility $\sum_{s \in \mathcal{S}} p(s) U(s, p')$ w.r.t. p', subject to the constraint $\sum_{s \in \mathcal{S}} p'_s = 1$; thence show that extremization occurs at $p' = p$.

[10] Everywhere strongly stable utility functions have been discussed in the statistics literature — outside the context of cognitive decision theory, but in a situation with identical mathematical desiderata — where such functions are known as 'proper scoring rules'. See, for example, Lad [3], Savage [7], for discussion and for lists of other such functions.

[11] Proof: Consider the totally peaked probability distributions, that assign probability 1 to some state $s \in S$ (and, of course, probability zero to all others $s' \neq s$). Relative to such a probability distribution, the expected epistemic utility $EU^p(p')$ of an arbitrary probability distribution p' just is the utility $U(s, p')$ of holding p' when state s obtains. But, if U is stable, then $EU^p(p')$ must be highest when $p' = p$. Thus, $U(s, \cdot)$ must be highest for the probability distribution p that is totally peaked on the state s. That is, if U is to be stable, U must encode the fact that the most epistemically favored probability distribution, when an arbitrary state s obtains, is the probability distribution that assigns credence 1 to the state s. That is, whatever the true state of the world, a maximum degree of belief in the true state is valued higher than any other credence distribution by any everywhere strongly stable utility function.

distribution near the true state than he would by having a 'flatter' probability distribution.

4.2 A somewhere stable utility function/Defense of cognitive decision theory

The following somewhere stable epistemic utility function has been discussed by Horwich [2]:127-9, Maher [5]:177-9, and Weintraub [12]:

$$\textbf{Linear utility function}: U_L(s,p) = p(s).$$

This utility function has an appealing mathematical simplicity, but, as Maher and Weintraub emphasize, it leads to very odd results. Specifically, the only credence distributions that are self-recommending with respect to this utility function are the totally peaked credence distributions ($p(s) = 1$ for some $s \in S$), and the indifferent distribution ($p(s) = \frac{1}{n}$ for each of n states, $n \leq |\mathcal{S}|$; again we assume that $\mathcal{S}$ has finite cardinality). If one were to hold any other credence distribution, one would maximize EU by shifting to a credence distribution that is assigns credence 1 to some disjunction of states that one currently considers most likely.[12]

What are we to make of this utility function? Maher and Weintraub think that it contains the seeds of an argument by *reductio* against the proposition (CDT), which the approach of present paper has taken as a premise:

CDT The dynamics of rational credence-distribution updating can be captured by a cognitive decision theory that recommends maximization of expected epistemic utility.

Maher's argument, with which Weintraub agrees, can be reconstructed as follows.[13]

P1 There exist (in logical space) rational agents who hold somewhere stable utility functions (SSUFs).

P2 If CDT is true, then, for any agent who holds a SSUF, spontaneous shifts from one credence function to another are sometimes rational.

[12] The following feature of this utility function should also be noted: it is not possible, by conditionalising on any proposition, to move from a probability distribution that (according to U_L) is self-recommending to one that is not. We have not, in fact, been able to find any somewhere stable utility function that does not possess this feature. If it could be shown that (under independently motivated auxiliary constraints) none exists, this would obviously strengthen our result (cf. our comments at the end of section 3).

[13] Our discussion in section 3.1 suggests that the possibility of rational spontaneous shifts is not the right way to characterize what is odd about somewhere stable utility functions — the point is rather that, given (CDT), a somewhere stable utility function has the consequence that some credence functions are 'forbidden points', in the sense that the ideally rational agent cannot hold one of these credence functions for any finite interval of time. However, we won't press this point. Maher's and Weintraub's argument applies, in any case, to the possibility of everywhere stable utility functions that are not everywhere *strongly* stable.

C1 If CDT is true, then there exist (in logical space) rational agents for whom spontaneous shifts are sometimes rational (from P1, P2)

P3 Such spontaneous shifts of credence are necessarily irrational: that is, nowhere in logical space are there *rational* agents who might sometimes perform spontaneous shifts.

C2 CDT is not true. (from C1, P3)

Clearly, if this argument were sound, the central claim of this paper (that Bayesians can justify conditionalisation by an appeal to maximization of expected epistemic utility) would be utterly undermined.

The argument is *valid*. Our objection to it is that either P1 or P3 is false, although we are not committed to a view as to which. We choose to insist on the correctness of the cognitive decision-theoretic approach, and argue by dilemma. Either shifts could be rational, or they could not. If they could, P3 is false, and so the argument fails. If they could not, P1 is false, and so the argument fails. In other words, either somewhere stable utility functions are to be ruled out as irrational, or they and their consequences are acceptable; in neither case is the CDT programme itself impugned.

Maher is aware of the possibility of this response. Since he insists absolutely on the irrationality of shifts (a view with which, as we noted in section 3.3 above, we are not unsympathetic), he gives serious consideration only to the possibility of rejecting P1. Maher's objection to this move is that it is 'completely *ad hoc*', since such a constraint on utility function lacks 'any **prior** plausibility' (*ibid.*, p.179; our emphasis in boldface). Our disagreement with Maher is methodological: we don't see why *prior* plausibility (i.e. prior to thinking through the consequences of adopting a somewhere stable utility function) should be required for rejection of P1. In any case, the problem with rejecting CDT in response to this argument is that that rejection is no less *ad hoc*: we are left with no convincing explanation of why one should maximize expected utility when choosing whether or not to go for a swim and when choosing whether or not to accept the proposition that humans are descended from apes, but not when choosing which credence distribution to adopt.

Incidentally, we *would* have a paradox for the decision-theoretic approach if we thought both that somewhere stable utility functions were pathological *and in addition* that no everywhere stable utility functions existed. But, as we have illustrated by example above, this latter condition does not obtain.

5 Conclusion

We have modelled the Bayesian agent's choice of updating policy as a decision problem within a cognitive decision theory. By doing so, it is possible to provide a justification for conditionalisation that appeals directly to the idea that epistemic rationality consists in taking steps that can reasonably be expected to lead to epistemically good outcomes. The justification is that, under independently motivated constraints, conditionalisation maximizes expected epistemic utility.[14]

References

[1] I. J. Good. On the principle of total evidence. *British Journal for the Philosophy of Science*, 17(3):319–321, February 1967.

[2] Paul Horwich. *Probability and evidence*. Cambridge University Press, Cambridge, 1982.

[3] Frank Lad. *Operational subjective statistical methods*. John Wiley & Sons, 1996.

[4] Isaac Levi. *Gambling with truth*. Knopf, New York, 1967.

[5] Patrick Maher. *Betting on theories*. Cambridge University Press, 1993.

[6] Graham Oddie. Conditionalization, cogency, and cognitive value. *British Journal for the Philosophy of Science*, 48:533–541, 1997.

[7] Leonard J. Savage. Elicitation of personal probabilities and expectations. *Journal of the Americal Statistical Association*, 66(366):783–801, December 1971.

[8] Leonard J. Savage. *The foundations of statistics*. Dover, New York, 2nd edition, 1972.

[9] Paul Teller. Conditionalization, observation, and change of preference. In William Harper and C. A. Hooker, editors, *Foundations of probability theory, statistical inference, and statistical theories of science*. D. Reidel, Dordrecht, 1976.

[10] Bas C. van Fraassen. *Laws and Symmetry*. Oxford University Press, New York, 1989.

[11] Bas C. Van Fraassen. A new argument for conditionalization. *Topoi*, 18:93–96, 1999.

[12] Ruth Weintraub. A Bayesian paradox. *British Journal for the Philosophy of Science*, 52:51–66, 2001.

[13] Paul Williams. Bayesian conditionalisation and the principle of minimum information. *The British Journal for the Philosophy of Science*, 31(2):131–144, June 1980.

[14] The authors would like to thank Frank Arntzenius and Wayne Myrvold for useful discussions. Thanks also to Jeremy Butterfield, for detailed comments on an earlier draft of this paper.