The value of information and the epistemology of inquiry

Richard Pettigrew
Richard.Pettigrew@bristol.ac.uk

August 11, 2024

Contents

Acknowledgements I'm very grateful to Mona Simion, Chris Kelp, Lilith Mace, and Claire Field for organising the Formal meets Normal conference at the Cogito research centre in Glasgow in July 2023; I'm grateful to the audience there for their excellent questions. This conference prompted me to delve into the literature on inquiry, and the current essay grew out of the talk I gave there. I'm also particularly grateful to Julia Staffel for highlighting a mistake I made when interpreting her account of transitional attitudes, and for working with me to figure out exactly where we agree and where we disagree.

Version control These notes were originally posted under the title 'Taking a Good look at the norms of gathering and responding to evidence'.

Chapter 1

Introduction

In the recent philosophical literature on inquiry, epistemologists have pointed out that their subdiscipline has typically begun at the point at which you already have your evidence, and it has focussed on identifying the beliefs for which that evidence provides justification or which count as knowledge for someone with that evidence. However, this recent literature goes on to point out, we are not mere passive recipients of the evidence we have.[1] While some of it comes to us unbidden, as we walk along the street, go about our day's work, or chop vegetables for dinner, we often actively collect it. We often choose to put ourselves in positions from which we'll gather some pieces of evidence but not others. We'll move to a position from which we'll see or hear or smell how the world is in one respect but miss how it is in another, as when I choose to inquire into the weather by looking out the window rather than at the forecast on the television screen. We'll prod the world in one way to see how it responds but we won't prod it in another, as when an physicist designs their experiment in this way rather than that. And so on.

As many in the recent literature point out, this has long been recognised, but typically epistemologists have taken the norms that govern inquiry to be practical, not epistemic. We inquire in order to find out things that inform our practical decisions, and so the decision what to find out is governed by practical considerations, and epistemologists leave well alone. Or, even if we inquire in order to find out things without an eye to their practical benefits, the things we do in order to inquire are not the sorts of thing for which one might have epistemic reasons. The recent literature challenges these assumptions and, as a result, uncovers a rich range of questions about the epistemic norms of inquiry.

In this essay, I approach these questions from the so-called formal side of epistemology, and specifically the Bayesian approach that occupies a large part of that side. The starting point for this approach is the observation that

[1] A small sample of recent writings in the epistemology of inquiry: (Hookway, 2006; Friedman, 2020; Kelp, 2021; Thorstad, 2022; Simion, 2023; Flores & Woodard, 2023; Rosa, 2025; Willard-Kyle, forthcoming; Staffel, ta).

beliefs and desires come in different strengths—I believe I'm Scottish more strongly than I believe I'll live to 90; I want this orange more strongly than I want that apple. What's more, we can represent these different strengths numerically, or in some other mathematically tractable way—I'm 99% confident I'm Scottish, but only 10% confident I'll live to 90; the amount by which I prefer an orange to an apple is the same as the amount by which I prefer a blackberry to a gooseberry. When we represent them numerically, we typically call the degrees of belief *credences* and the degrees of desire *utilities*.

Among other things, Bayesians seek norms that govern how you should set your credences initially, how you should change them when new evidence comes in, and how you should combine them with your utilities, and possibly other attitudes, in order to set the preferences that guide your actions. The non-formal side of epistemology, on the other hand, has traditionally been more concerned with full beliefs—ones that don't come in degrees. Among other things, they have been interested in the conditions under which these beliefs have certain features: when they are justified; when they amount to knowledge; when they have the authority to guide action, ground assertion, or figure in further deliberation.

I'd like to persuade you that it's natural to look to the Bayesian approach to help us think about inquiry. For one thing, whether to inquire or not is a choice, as is the decision to inquire in this way rather than that, and Bayesian epistemology offers a number of sophisticated theories of rational choice and decision-making. For another, while inquiring can lead us to adopt full beliefs or abandon them, it can often change our attitudes in more fine-grained ways that leave our full beliefs as they were.[2] Perhaps I am 50% confident in something and my inquiry might make me either 49% confident in it or 51% confident in it, depending on which way the evidence that I collect goes. In this case, I don't have a full belief to begin with, and neither way my evidence goes will lead me to have one afterwards; and yet there is real change, and we must be able to represent that change if we're to understand inquiry, for often the way we come to change our full beliefs is by the gradual accretion of evidence that eventually leads us to believe or disbelieve or suspend judgment.

What's more, because current Bayesian epistemology developed out of the philosophy of science, where it was used to understand how certain sorts of evidence support certain sorts of conclusion, this branch of epistemology has long discussed inquiry. The standard reference in the Bayesian philosophy of science is I. J. Good's (1967) 'On the Principle of Total Evidence' or Leonard J. Savage's (1954) *The Foundations of Statistics*. However, the result at the centre of Good's argument is already proved by Janina Hosiasson in her 1931 *Mind* paper, 'Why do we prefer probabilities relative to many data?'.[3] We know

[2] Arianna Falbo (2023) makes a related point.

[3] I learned of Hosiasson's paper from Christian Torsell's (2024) paper 'Janina Hosiasson and the Value of Evidence', and wholeheartedly recommend those interested in the history of this tradition to consult that. I will refer to her as Janina Hosiasson throughout, since that is the name under which she published this crucial paper. However, she later became Janina Hosiasson-Lindenbaum, and her more influential papers in the then-nascent field of Bayesian philosophy of

from his notes that Frank P. Ramsey (1990) was also aware of it, though he never published it. And, in economics, the standard reference is David Blackwell's (1951) 'Comparison of Experiments', which extends the result considerably in a direction I won't pursue in this essay. All use the same account of the pragmatic utility of gathering evidence, which originates in Hosiasson's paper; this is the basis for all of the future results.

Building on these insights, Bayesian philosophy of science and its more recent expansion into Bayesian epistemology more broadly has produced a reasonably well-developed framework in which to understand norms of inquiry, both epistemic and practical. The first half of this essay presents this framework. In Chapter 2, I will present the framework that has been built on Hosiasson's insight. Its centrepiece is the pragmatic version of the so-called Value of Information Theorem, which formalizes the intuitive idea that you're better off gathering as much information as you can before making a decision, providing the process of gathering it isn't too costly. So, for instance, it's better for me to check the weather forecast before I decide whether or not to take my umbrella with me when I leave the house, providing it doesn't take me too much time. More precisely, it says that, if the information available is free, you'll never do worse in expectation by taking it before making the decision, and you'll very often do better. This tells us something about the pragmatic value of inquiry: it aids decision-making. I'll present the theorem, as well as a series of recent generalizations due to John Geanakoplos (1989), Nilanjan Das (2023), and Kevin Dorst et al. (2021).

I'll also note how things go differently if you use a decision theory other than expected utility theory, such as Lara Buchak's (2013) risk-weighted expected utility or Chris Bottomley and Timothy Luke Williamson's (ta) weighted-linear utility theory; and also how things go if you have imprecise probabilities instead of precise ones.

In Chapter 3, I will present a purely epistemic version of the Value of Information Theorem, due to Graham Oddie (1997). This theorem formalizes the intuitive idea that you'll have better beliefs the more you inquire. So, for instance, my credences about whether it will rain or not will likely be better if I check the weather forecast than if I don't. More precisely, it says that gathering evidence never decreases the accuracy of your credences in expectation, and very often it increases it. This tells us something about the epistemic value of inquiry: it gets us better beliefs. Again, I'll present the theorem, as well as a series of recent generalizations due to Wayne Myrvold (2012), Alejandro Pérez Carballo (2018), and, again, Kevin Dorst et al. (2021).

In the second half of this essay, I put to work the approach to inquiry I've been describing in the first half. I turn to some of the questions from the recent debate about inquiry and ask how the approach initiated by Hosiasson and adapted by Oddie can help us answer them. Questions will include: When should we initiate an inquiry, when should we continue it, when should we conclude it, and when should we reopen it (Chapter 4)? How do epistemic

science were published under that name.

norms of inquiry relate to epistemic norms of belief or credence, and can they conflict, as Jane Friedman (2020) contends? How should we resolve the apparent puzzle raised by Friedman's example of counting the windows in the Chrysler Building? How should we direct our attention, as Georgi Gardiner (2022) asks (Chapter 5)? How should we choose and reason after receiving evidence we'd rather not have received (Chapter 6)? And how should we understand the epistemic error that occurs when someone is resistant to evidence in the way Mona Simion (2023) describes (Chapter 7)?

Throughout, I will present the ideas both informally and formally. I'll place any formal presentation that uses mathematical notation in blue boxes. These can be skipped over, if you prefer, as the ideas are presented informally in the surrounding text. But I include them, as some people will find it easier to see the ideas presented in that notation, and hopefully some will wish to use this framework and these results themselves, and for that you'll want to be sure that you're formalising things as others have.

Before we begin, I should note that there is another facet to the study of inquiry in formal epistemology: it is the study of collective rather than individual inquiry, and it tends to ask how we should structure our scientific communities, institutions, and practices in order to best discover the truth together.[4] I will not discuss it here, but only because it makes less obvious contact with the questions raised in the recent work on inquiry in non-formal epistemology; though of course it makes important contact with other work in non-formal epistemology, namely, on echo chambers, epistemic bubbles, misinformation, conspiracy theories, and testimony.

[4]See, for instance, (Zollman, 2007, 2010; Rosenstock et al., 2017; Bright, 2017; O'Connor & Weatherall, 2018).

Part I

The value of information

Chapter 2

The pragmatic value of information

While most discussion of the value of evidence or information focuses on what has become known as the Value of Information Theorem, the real contribution lies in the definition of the pragmatic value of gathering evidence, which Ramsey, Hosiasson, Blackwell, and Good all give, however independently. In this chapter, I'll introduce the framework of credences in a little more detail (Section 2.1), give this definition of the pragmatic value of gathering evidence (Section 2.2), and state the Value of Information Theorem (Section 2.3).

2.1 Credences

Hosiasson's definition and the Pragmatic Value of Information Theorem were stated and proved in the standard Bayesian framework, and in this framework we represent someone's beliefs by their *credences*. These are the states we ascribe when we say things like 'Ada is 65% sure it's going to rain' or 'Cal is 50-50 whether they left on the gas'. In these cases, we say Ada has credence 0.65 in the proposition that it's going to rain, and Cal has credence 0.5 in the proposition they left on the gas. Sometimes, these are known as *degrees of belief*, or *strengths of confidence*, *partial beliefs*, or *personal* or *subjective probabilities*. I'll talk of credences throughout.

An individual's credences are numbers assigned to propositions to measure the strength of the individual's beliefs in those propositions. So, what are the propositions? We'll take a pretty flat-footed approach here.[1] We'll assume there's a set of possible worlds that represent the possibilities the individual entertains at the finest level of grain at which they entertain them. And then we'll represent propositions as sets of these possible worlds. And we'll say that the individual assigns credences to propositions thus represented. So, for instance, an individual might only distinguish three states of the world:

- *The Eiffel Tower is taller than the Taj Mahal* (w_1);

[1]Taking this approach is not essential to the framework and results I'm going to present; it'll just make it easier. We could just as well take the objects of credences to be sentences or propositions represented in some other way than as sets of possible worlds.

- *The Eiffel Tower is the same height as the Taj Majal* (w_2);
- *The Taj Mahal is taller than the Eiffel Tower* (w_3).

And then the propositions to which they'll assign credences are those represented by the following sets:

- the empty set ($\emptyset$), which represents the proposition that is true at no world—the necessarily false proposition, if you like;
- the singleton of each of the three states, which represents the proposition true at that state and no other ($\{w_1\}$, $\{w_2\}$, $\{w_3\}$);
- the three pairs from the states, which represent *The Eiffel Tower is at least as tall as the Taj Mahal* ($\{w_1, w_2\}$), *The Taj Mahal is at least as tall as the Eiffel Tower* ($\{w_2, w_3\}$), and *The Taj Mahal and the Eiffel Tower are different heights* ($\{w_1, w_3\}$); and
- the set of all three states, which is true at each of them ($\{w_1, w_2, w_3\}$)—the necessarily true proposition.

We call the set of propositions to which the individual assigns credences their *agenda*. For ease of exposition, we assume that there are only finitely many possible worlds they entertain, and so their agenda is similarly finite.

Our individual's *credence function* is then the function that takes each proposition in their agenda and assigns to it the number at least 0 and and at most 1 that represents their credence in it; the number that measures the strength of their belief in it. 0 represents minimal credence or least possible strength of belief; 1 represents maximal credence or greatest possible strength of belief.

We'll assume throughout that our inquiring individual's credence function obeys the Bayesian norm of Probabilism, which says that their credence function must be a probability function. That is, it must assign credence 1 to the necessary truth, namely, the proposition represented by the set of all possibilities; it must assign credence 0 to the necessary falsehood, namely, the proposition represented by the empty set; and the credence it assigns to a disjunction of two mutually exclusive propositions must be the sum of the credences it assigns to the disjuncts.

Suppose W is a finite set of possible worlds, and let $\mathcal{F}$ be the set of subsets of W. Then a *credence function* is a function $C : \mathcal{F} \to [0, 1]$. And a credence function is *probabilistic* iff

(i) $C(W) = 1$;

(ii) $C(\emptyset) = 0$;

(iii) $C(X \cup Y) = C(X) + C(Y)$ if $X \cap Y = \emptyset$.

Equivalently,

(i) $\sum_{w \in W} C(w) = 1$;

(ii) $C(X) = \sum_{w \in X} C(w)$.

where we abuse notation and write $C(w)$ for the credence that C assigns to the singleton set $\{w\}$.

Probabilism Rationality requires that your credence function is probabilistic.

2.2 The pragmatic value of gathering evidence

As I mentioned above, the key insight in the Bayesian approach to inquiry is the definition of the pragmatic value of gathering evidence, which was given originally by Janina Hosiasson, Frank Ramsey, David Blackwell, and I. J. Good at different times. Good's is most commonly discussed in philosophy of science, Blackwell's in economics, and Hosiasson's version has been sadly neglected, despite being published first.

This definition begins with another definition; it begins with a definition of the pragmatic value of a probabilistic credence function relative to a decision you will face. Suppose you will face a particular decision between a range of options, where an option is specified by giving its utility at each possible state of the world, and the utility of an option at a world is a real number that measures how much you value the outcome of that option at that world. Then the standard theory of choice under uncertainty says that you should pick an option with maximal expected utility from the point of view of the credence function you have when you face the decision: that is, you calculate the expected utility of each option by taking its utility at each world, weighting that by the credence you assign to that world, and summing up these credence-weighted utilities; and then you should pick an option whose expected utility is maximal—that is, no other option has higher expected utility, though perhaps others have equally high expected utility.

So let's assume you'll do this. Then we can define the pragmatic value for you, at a particular state of the world, of having a particular credence function when faced with a particular decision: it is the utility, at that state of the world, of the option that this credence function will lead you to pick from those available in the decision. This will be one of the options that maximizes expected utility from the point of view of that credence function; and since there might be more than one that maximizes that, we must assume you have a way of breaking ties between them.

The pragmatic utility of a credence function

Some preliminary definitions:[a]

- A *decision problem* $\mathcal{D}$ is a set of options.
- Each *option* o in $\mathcal{D}$ is a function from the set of possible worlds W into the real numbers: $o(w)$ is the utility of o at w.
- Given a probabilistic credence function C, the *expected utility of o from the point of view of C* is $\sum_{w \in W} C(w)o(w)$.
- Given a decision problem $\mathcal{D}$ and a probabilistic credence function C, let $\mathcal{D}_C$ be the *choice set*: that is, it is the set of options in $\mathcal{D}$ that maximize expected utility from the point of view of C.
- A *tie-breaker function* τ takes any set of options and picks one of them out. Our individual uses it when there is more than one option that maximizes expected utility. Given a decision problem $\mathcal{D}$ and a credence function C, they apply τ to the choice set $\mathcal{D}_C$ to give the option $\tau(\mathcal{D}_C)$ that they pick.

Definition 1 (Pragmatic utility of a credence function)**.** *The pragmatic utility, at world w, of a credence function C relative to decision problem $\mathcal{D}$ and tie-breaker function τ, is*

$$\mathrm{PU}^{\mathcal{D},\tau}(C, w) = \tau(\mathcal{D}_C)(w)$$

That is, it is the utility, at w, of $\tau(\mathcal{D}_C)$, which is the option our individual will pick from among those options in $\mathcal{D}$ that maximize expected utility relative to their credence function C.

[a] For our initial presentation of the Value of Information framework, we work within Savage's (1954) version of expected utility theory. In Section 2.4, we consider the status of the Value of Information Theorem in other versions.

So, for instance, suppose I have to walk to the shops and I must decide whether or not to take an umbrella with me. And suppose I have credences concerning whether or not it will rain as I walk there. Let's suppose first that taking the umbrella uniquely maximizes expected value from the point of view of those credences. Then the pragmatic value of those credences at a world at which it does rain is the utility of walking to the shops in the rain with an umbrella, while their pragmatic value at a world at which it doesn't rain is the utility of walking to the shops with no rain carrying an umbrella. And similarly, if leaving without the umbrella uniquely maximizes expected utility from the point of view of those credences, then their pragmatic value at a rainy world is the utility of walking to the shops in the rain without an umbrella, and their pragmatic value at a dry world is the utility of walking to the shops with no rain and no umbrella. And if they both maximize expected utility from the point of view of the credences, then the pragmatic value of the credences will depend on how I break ties. So, holding fixed the decision problem you'll face and the way you break ties, the pragmatic value of a credence function is the utility of the option it'll lead you to pick.

Now, having defined the pragmatic value of a credence function relative to a particular decision you'll face and a way of breaking ties, we can define the pragmatic value of a particular episode of evidence-gathering relative to such a decision and tie-breaker function. We represent such an episode as follows: for each possible state of the world, we specify the strongest proposition you'll learn as evidence at that state of the world—this is what Nilanjan Das (2023) calls an *evidence function*. And we assume that you have a plan for how to respond to each possible piece of evidence—we call this an *updating plan*. Then the pragmatic value, at a particular world, of an episode of evidence-gathering is the pragmatic value, at that world, of the credence function you'll have after learning whatever evidence you'll gather at that world and responding to it in the way your updating plan says you should. So, holding fixed the decision problem you'll face and the way you break ties, the pragmatic value of a credence function is the utility of the option it'll lead you to pick, and the pragmatic value of gathering evidence is the pragmatic value of the credence function it will lead you to have when you respond to that evidence as you plan to.

Of course, this assumes that it is already fixed how you will respond to any evidence you receive; and indeed Good assumes you'll update as the Bayesian says you should: you'll condition on whatever proposition you receive as evidence; that is, your unconditional credence in a proposition after receiving some evidence is your prior conditional credence in that proposition given the evidence you learn, so that my posterior credence in rain after learning the forecast is dry is my prior conditional credence in rain given the forecast says dry. For the moment, we'll stick with this assumption; later, we'll lift it to see what happens.

The pragmatic utility of an evidence-gathering episode

Some preliminary definitions:

- We represent an evidence-gathering episode by an *evidence function* $\mathcal{E} : W \to \mathcal{F}$. This takes each world w to the strongest proposition $\mathcal{E}_w$ you learn in that world if you gather the evidence.
- Given an evidence-gathering episode $\mathcal{E}$ and a prior credence function C, if you engage in the evidence-gathering episode with that prior, your posterior at world w will be $C(- \mid \mathcal{E}_w)$, providing $C(\mathcal{E}_w) > 0$.[a]

Definition 2 (Pragmatic utility of an evidence-gathering episode)**.** *The pragmatic utility, at world w, of an evidence-gathering episode $\mathcal{E}$, relative to decision problem $\mathcal{D}$ and tie-breaker τ, is*

$$\mathrm{PU}^{\mathcal{D},\tau}(\mathcal{E}, w) = \mathrm{PU}^{\mathcal{D},\tau}(C(- \mid \mathcal{E}_w), w).$$

That is, it is the utility, at w, of the posterior credence function $C(- \mid \mathcal{E}_w)$ that you will have after learning the evidence you'll learn at that world.

[a] Where $C(X \mid Y) = C(XY)/C(Y)$, when $C(Y) > 0$.

So, for instance, suppose I have to walk to the shops later and, at that point, I'll have to decide whether or not to take an umbrella with me. And suppose that, between now and then, I can gather evidence by looking at the weather forecast. If I do, I'll learn one of two things: rain is forecast, or rain is not forecast. And updating on that evidence as I plan to, should I choose to gather it, might well change my credences concerning whether or not it will rain on my way to the shops. Then what is the value, at a particular state of the world, of gathering evidence by looking at the forecast? Consider a world at which (i) rain is not forecast but (ii) it does rain; and suppose that, upon learning that rain is not forecast, I'll drop my credence in rain low enough that I'll not take my umbrella. Then the value of gathering evidence at that world is the utility of walking to the shops in the rain without an umbrella. In contrast, consider a world at which (i) rain is forecast but (ii) it doesn't rain; and suppose that, upon learning that rain is forecast, I raise my credence in rain high enough that I take the umbrella. Then the value of gathering evidence at that world is the utility of walking to the shops with no rain but carrying an umbrella. And so on.

This is the Hosiasson-Blackwell-Good account of the pragmatic value, at a particular world, of a particular episode of evidence-gathering; and it is relative to the decision problem you'll face with the credences you come to have after updating, and the way you'll break ties between the options, if you need to. With this in hand, we can now define the *expected* pragmatic value of such an episode from the point of view of your prior credence function (or indeed from the point of view of any probability function). And we can also define the expected pragmatic value of not gathering evidence at all, since that is just the degenerate case of evidence-gathering in which you simply learn a tautology at every state of the world.

The expected pragmatic utility of an evidence-gathering episode

The expected utility of an evidence-gathering episode $\mathcal{E}$, from the point of view of C and relative to decision problem $\mathcal{D}$ and tie-breaking function τ, is

$$\text{Exp}_C(\text{PU}^{\mathcal{D},\tau}(\mathcal{E})) = \sum_{w \in W} C(w)\text{PU}^{\mathcal{D},\tau}(\mathcal{E}, w) = \sum_{w \in W} C(w)\tau(\mathcal{D}_{C(-|\mathcal{E}_w)})(w).$$

The expected utility of not gathering evidence, from the point of view of C and relative to decision problem $\mathcal{D}$ and tie-breaking function τ, is

$$\text{Exp}_C(\text{PU}^{\mathcal{D},\tau}(C)) = \sum_{w \in W} C(w)\text{PU}^{\mathcal{D},\tau}(C, w) = \sum_{w \in W} C(w)\tau(\mathcal{D}_C)(w).$$

Huttegger and Nielsen's alternative definition

It is worth noting here, before we move on, that, following Brian Skyrms (1982, 1990b), Simon Huttegger and Michael Nielsen use a different definition of the expected pragmatic utility of an evidence-gathering episode (Huttegger, 2013, 2014; Huttegger & Nielsen, 2020; Nielsen, 2024). It is equivalent to the one I have given when $\mathcal{E}$ is factive and partitional, but not otherwise. I'll state it here and give an example that motivates using the definition I use instead.

The Skyrms-Huttegger-Nielsen expected utility of an evidence-gathering episode $\mathcal{E}$, from the point of view of C and relative to $\mathcal{D}$ and τ, is

$$\sum_{w \in W} C(w) \sum_{w' \in W} C(w' \mid \mathcal{E}_w)\tau(\mathcal{D}_{C(-|\mathcal{E}_w)})(w')$$

That is, it is the prior's expectation of the posterior's expectations of the choice the posterior will make. In contrast, the Hosiasson-Blackwell-Good definition takes it to be the prior's expectation of the choice the posterior will make.

I favour the Hosiasson-Blackwell-Good definition given above because I think the Skyrms-Huttegger-Nielsen definition gives the wrong answer in the following case, which is essentially a simplified version of Tim Williamson's unmarked clock case, which we'll meet below:[a]

Example 1. *Suppose we are in the following situation:*

- *The possible worlds:* $W = \{1, 2, 3, 4\}$.
- *Your prior credence function:* $C(1) = C(2) = C(3) = C(4) = 1/4$.
- *The evidence function:*
  - $\mathcal{E}_1 = 4 \vee 1 \vee 2$;
  - $\mathcal{E}_2 = 1 \vee 2 \vee 3$;
  - $\mathcal{E}_3 = 2 \vee 3 \vee 4$;
  - $\mathcal{E}_4 = 3 \vee 4 \vee 1$.
- *You'll update by conditionalizing on whatever evidence you receive. So your posterior at world 2, for instance, will be* $C(- \mid \mathcal{E}_2) = C(- \mid 1 \vee 2 \vee 3)$.

Now consider the decision between the following options:

	1	2	3	4
Odd	4	−6	4	−6
Even	−6	4	−6	4
Reject	0	0	0	0

Essentially, Odd is a 10 utile bet on the world being odd priced at 6 utiles, Even is a 10 utile bet on the world being even priced at 6 utiles, and Reject is the rejection of both bets at that price.

Then your priors will pick Reject. If you're in an odd-numbered world, your posteriors will pick Even, and they'll therefore lose 6 utiles. If you're in an even-numbered world, your posteriors will pick Odd, and they'll therefore lose 6 utiles. According to the Hosiasson-Blackwell-Good definition of the value of information, learning this evidence has lower expected value than not learning, because you know that learning will lead you to lose 6 utiles, while sticking with your priors will lead you to lose no utiles. But according to the Skyrms-Huttegger-Nielsen definition, learning has greater expected value than not learning, because each posterior assigns positive expected value to the option it picks, while the prior assigns expected value zero to the option it picks. I favour the Hosiasson-Blackwell-Good definition as it gives what seems to me the right answer here.

[a] As Michael Nielsen pointed out to me, Dmitri Gallow (2021, Section 3.2, Appendix A) gives a similar sort of case in a slightly different context.

We are now furnished with all the definitions we need to state Good's Value of Information Theorem.

2.3 The Value of Information Theorem

Good's Value of Information Theorem runs as follows: Fix a decision problem you'll face at a later time; fix the way you'll break ties between a set of options when they all maximize expected utility; and assume you plan to update upon receiving evidence in the way the Bayesian suggests, namely, by conditionalizing on it. Now suppose that, for no cost, you may gather some evidence before you face the decision problem. And suppose that the evidence function that represents this possible evidence-gathering episode is *factive* and *partitional*: to say that it is factive is to say the proposition you'll learn is true; to say it's partitional is to say that the set of propositions you might learn forms a partition—for any possible world, exactly one of these propositions is true at that world. Then the expected pragmatic value, from the point of view of your prior credences, of gathering that evidence is at least as great as the expected pragmatic value, again from the point of view of your prior credences, of not gathering it; and, very often, it is strictly greater. How often? Well, if you assign some positive credence to getting evidence that will result in posterior credences from whose point of view the option that your priors would have chosen is no longer optimal, then the expected pragmatic value of gathering the evidence is strictly greater than the expected pragmatic value of not gathering it. In slogan form: it's always rationally permissible to take free evidence, and it's rationally required when you think it might lead you to consider what you would have chosen without it suboptimal.

Hosiasson's Pragmatic Value of Information Theorem

Some preliminary definitions: Suppose $\mathcal{E}$ is an evidence function. Then:

- $\mathcal{E}$ is *factive* if, for each world w in W, $\mathcal{E}_w$ is true at w;
- $\mathcal{E}$ is *partitional* if $\{E_w : w \in W\}$ forms a partition.

Theorem 1 (The Value of Information Theorem)**.** *If $\mathcal{E}$ is factive and partitional, and C is a probabilistic credence function, then*

(i)
$$\text{Exp}_C(\text{PU}^{\mathcal{D},\tau}(\mathcal{E})) \geq \text{Exp}_C(\text{PU}^{\mathcal{D},\tau}(C))$$

(ii)
$$\text{Exp}_C(\text{PU}^{\mathcal{D},\tau}(\mathcal{E})) > \text{Exp}_C(\text{PU}^{\mathcal{D},\tau}(C))$$

if there is w in W such that $C(w) > 0$ and $\tau(\mathcal{D}_C) \notin \mathcal{D}_{C(-|\mathcal{E}_w)}$.

In fact, Hosiasson proves a more general result from which the Value of Information theorem follows. To state it, we need a definition:

- *$\mathcal{E}$ is at least as informative as $\mathcal{E}'$* if, for any world w, $\mathcal{E}_w \subseteq \mathcal{E}'_w$.

Lemma 2 (The Value of Information Lemma)**.** *If $\mathcal{E}$ and $\mathcal{E}'$ are factive and partitional, C is a probabilistic credence function, and $\mathcal{E}$ is at least as informative as $\mathcal{E}'$, then*

(i)
$$\text{Exp}_C(\text{PU}^{\mathcal{D},\tau}(\mathcal{E})) \geq \text{Exp}_C(\text{PU}^{\mathcal{D},\tau}(\mathcal{E}'))$$

(ii)
$$\text{Exp}_C(\text{PU}^{\mathcal{D},\tau}(\mathcal{E})) > \text{Exp}_C(\text{PU}^{\mathcal{D},\tau}(\mathcal{E}'))$$

if there is w in W such that $C(w) > 0$ and $\tau(\mathcal{D}_{\mathcal{E}'_w}) \notin \mathcal{D}_{C(-|\mathcal{E}_w)}$.

Theorem 1 follows from Lemma 2 if we note that not gathering evidence is the same as engaging in the degenerate evidence-gathering episode $\mathcal{E}$, where $\mathcal{E}_w$ is the necessary proposition for all w in W, and any factive and partitional evidence function is at least as informative as this degenerate one.

A natural response upon first encountering the Value of Information Theorem is to think it's obviously true. After all, surely any true evidence is guaranteed to improve your epistemic situation, and surely improving the epistemic standpoint from which you face decisions leads to better choices. However, this isn't true. Evidence can be true but misleading. It is easy to find examples in which you would have made a decision that obtained for you more utility had you not learned the true evidence you did before choosing.

Suppose you know it's either sunny, rainy, or windy outside, and you divide your credences equally over the three. In fact, it's windy. You can stay indoors, or you can go outside. Staying indoors gets you 8 units of utility for sure; if you go outside and it's sunny, you get 14 units, if it's windy you get 6, and if it's rainy, you get 1.

	Sunny	Windy	Rainy
Indoors	8	8	8
Outdoors	14	6	1

Then you currently prefer to stay indoors, since the expected utility of doing that (8) is higher than the expected utility of going outside ($\frac{14+6+1}{3} = 7$). Now you learn it's sunny or windy. You update on this information and come to prefer going outside, since your new expected utility for doing that ($\frac{14+6}{2} = 10$) is higher than your new expected utility for staying indoors (which is still 8). But, since it's windy, you end up less well off as a result of learning then choosing; you'd have done better just to choose without learning.

So the Value of Information theorem is not obviously true. So why does it hold? If your evidence will teach you which member of a pre-specified partition is true, while misleading evidence is possible, when its effects are weighted by your credence you'll get it and considered together with the possibility of non-misleading evidence, whose effects are weighted by your credence you'll get them, it turns out that the possibility of non-misleading evidence wins out and it's better, in expectation, to gather the evidence and take the risk. This is helpful to bear in mind as we consider versions of the theorem in other settings in Sections 2.5 and 2.6.

2.4 The Sure Thing Principle

A natural way to see that the Value of Information theorem is true begins with Leonard Savage's *Sure Thing Principle*. This is a putative norm of decision theory—indeed, one of the axioms Savage laid done for rational decision-making. On one reading, it tells us how our preferences after learning something should hang together with our preferences beforehand. Here is the story with which Savage introduces it:

> A businessman contemplates buying a certain piece of property. He considers the outcome of the next presidential election relevant. So, to clarify the matter to himself, he asks whether he would

> buy if he knew that the Democratic candidate were going to win, and decides that he would. Similarly, he considers whether he would buy if he knew that the Republican candidate were going to win, and again finds that he would. Seeing that he would buy in either event, he decides that he should buy, even though he does not know which event obtains, or will obtain, as we would ordinarily say. (Savage, 1954, 21)

The general principle is this: If you would prefer one option (buying the house) to another (not buying it) were you to learn that a given proposition is true (the Democrat wins), and you would prefer that first option to that second option were you to learn that same proposition is false (the Republican wins), then you should prefer the first option to the second before learning whether it's true or false—here, as before, we assume that, if you learn a proposition, you update your credences by conditioning on it in the way the Bayesian requires you to do.

Standard expected utility theory, which tells you to set your preferences in line with your expected utilities, satisfies this principle, and we can use that fact to show that the Value of Information Theorem is true. Suppose you have the opportunity to learn whether a proposition is true or false before you make a decision. Suppose you learn it's true. Will you be glad you did? That is, will you think the decision you'll now make is at least as good as the decision you would have made had you not learned the proposition? Yes! For you're asking whether your current credences expect the option you'll now choose with those current credences to have at least as great expected utility as the option you would have chosen with your old credences. And of course the answer is that they will. What's more, if your old credences would have led you to choose differently, then your current credences expect the choice you make with them to be strictly better! And the same goes if you were to learn the proposition is false. The credences you'd then have after updating on that information would expect the option you'd choose with them to have maximal expected utility, and be strictly better utility if your old credences would have led you to choose differently. And so, either way, you'd be glad you learned the evidence: and so, by the Sure Thing Principle, prior to learning it, you should prefer learning it to not learning it.

The Sure Thing Principle and Expected Utility Theory

While Savage's example and my discussion of it have stated the Sure Thing Principle informally in terms of what you'd choose now and what you'd choose were you to learn something, the standard statement of principle in fact involves the relationship between your unconditional preference ordering over options and, for each proposition, a conditional preference ordering over options under the supposition of that proposition. This is related to the version concerning learning that

I have been discussing informally if we assume that, when we learn a proposition, our new unconditional preferences should be our old conditional preferences under the supposition of that proposition.

We'll write the unconditional preference ordering $\preceq$ and, for each E, the conditional preference ordering $\preceq_E$.

> **The Sure Thing Principle** If $o_1 \preceq_E o_2$ and $o_1 \preceq_{\overline{E}} o_2$, then $o_1 \preceq o_2$.

In expected utility theory, we say that an individual's unconditional preferences order the options by their expected utility relative to their credence function, so that

$$o_1 \preceq o_2 \Leftrightarrow \mathrm{Exp}_C(o_1) \leq \mathrm{Exp}_C(o_2)$$

And we say that their conditional preferences order the options by their expected utility relative to their conditional credence function, so that

$$o_1 \preceq_E o_2 \Leftrightarrow \mathrm{Exp}_{C(-|E)}(o_1) \leq \mathrm{Exp}_{C(-|E)}(o_2)$$

To see that expected utility theory satisfies the Sure Thing Principle, we note the following fact about expected utility theory:

$$\mathrm{Exp}_C(o) = C(E)\mathrm{Exp}_{C(-|E)}(o) + C(\overline{E})\mathrm{Exp}_{C(-|\overline{E})}(o)$$

That is, the expectation of an option is the expectation of its conditional expectations. So if, for both E and $\overline{E}$, the conditional expectation of o_1 is at most the conditional expectation of o_2, then the expectation of the conditional expectations of o_1 must be at most the expectation of the conditional expectations of o_2. And so the expectation of o_1 must be at most the expectation of o_2.

2.5 The value of information in other theories of rational choice

So far, I have been assuming that the correct theory of rational choice is expected utility theory, and indeed I've been using Savage's version of it. But what do alternative decision theories say about the value of gathering evidence?

Suppositional decision theories

I'll begin with the so-called *suppositional decision theories* and describe Michael Nielsen's (2024) elegant result that any such theory that satisfies the value of information must be a particular sort of causal decision theory, namely, one based on what David Lewis (1981) called an *imaging operators*.

In Savage's decision theory, the options between which we choose are represented by functions from worlds to utilities, and we define the value of an option to be its expected utility, that is, the sum over each possible world of the utility of that option at that world weighted by the credence assigned to the world. Notice that, in this theory, the credence assigned to a world does not depend on the option whose value we are assessing. Savage's theory thereby assumes what is sometimes called *act-state independence*: performing the action does not affect the state of the world.

In suppositional theories, we do not assume this. A suppositional theory is defined by a supposition function, which takes the decision-maker's credence function and an option and returns a new credence function, which we think of as giving the decision-maker's credences *under the supposition* that the option is chosen. Then the value of an option according to this theory is just the expected utility of performing that option under the supposition that it is performed, that is, the sum over each possible world of the utility of that world weighted by the credence in that world under the supposition that the option is chosen. Notice that, in such theories, worlds have the same utility, regardless of which option is chosen. In Savage's theory, in contrast, the utility of the world depends on which option is chosen. In suppositional theories, instead, the option does not alter the utility of the world; it alters the credence in the world that is used to weight the fixed utility of that world.

There are many suppositional theories out there. Richard Jeffrey's (1983) evidential decision theory is a suppositional theory where the credence in a world under the supposition of choosing an option is just the conditional credence in that world given that the option is chosen. Allan Gibbard and William L. Harper's (1978) version of causal decision theory is a suppositional theory where the credence in a given world under the supposition of choosing an option is the probability of the counterfactual that says that if that option were chosen, the world would be that one. And there are others.

Nielsen is particularly interested in the suppositional theories generated by so-called imaging operators. An imaging operator is generated by a *selection function*. A selection function takes an option o and a world w and says, for each possible world w', what proportion of the original credence function's probability in w' should be transferred to w under the supposition that o is performed (a proposition we will write as $\overline{o}$). So, under that supposition, the imaging operator's credence in a world w is the sum, over each possible world w', of the original credence in w', weighted by the selection function's value for the w' under the supposition $\overline{o}$. So the imaging function works through each possible world w', taking a little of its credence to give to w, where the proportion it gives is specified by the selection function.

Nielsen then proves that the Value of Information Theorem holds for a suppositional theory—that is, for every decision problem and every finite partition, the theory values at least as much learning the true element of the partition and then deciding as it does deciding without learning and sometimes more—if, and only if, it is generated by an imaging operator—that is, the supposition function is an imaging operator.

Indeed, Nielsen's result tells us something stronger than that. For he doesn't assume upfront that you learn by conditionalizing on the evidence; rather, he proves that learning by the evidence and choosing in line with a suppositional theory with a supposition function that is an imaging operator is necessary and sufficient for the Value of Information Theorem.

One upshot of Nielsen's result, which we already knew from Brian Skyrms (1990b) is that the Value of Information theorem does not hold of Jeffrey's evidential decision theory, since that cannot be generated by a supposition function generated by an imaging operator.

Suppositional decision theories and Nielsen's theorem

- A *supposition function* s takes a credence function C and a proposition A and returns a credence function $s(C, A)$, where, if $C(w)$ gives the credence in w under no supposition, then $s(C, A)(w)$ gives the credence in world w *under the supposition of* A. We assume $s(C, A)(A) = 1$.
- Given an option o, we write $\overline{o}$ for the proposition that o is performed.
- A decision theory is *suppositional* if it says that you should choose an option with maximal suppositional expected utility relative to a supposition function, where the suppositional expected utility of o relative to supposition function s is given by

$$\sum_{w \in W} s(C, \overline{o})(w) U(w)$$

where $U(w)$ is the utility of world w.

- Some suppositional theories:

(i) Evidential decision theory:

$$s(C, A)(w) = C(w \mid A).$$

(ii) Counterfactual causal decision theory:

$$s(C, A)(w) = C(A \,\Box\!\!\rightarrow w).$$

(iii) Stalnakerian causal decision theory:

$$s(C, A)(w) = \sum_{w' \in W} C(w') f(A, w')(w),$$

where f is an *imaging operator*, that is, $f(A, w')$ is a probability function for each option A and world w', and $f(A, w')(A) = 1$.

- A learning operator l takes a credence function C and a proposition E and returns a credence function $l(C, E)$, where, if $C(w)$ gives the prior credence in w, then $l(C, E)(w)$ gives the posterior credence in world w *upon learning* E. We assume $l(C, E)(E) = 1$.

Theorem 3 ((Nielsen, 2024)). *The following are equivalent:*

(i) *For every credence function C, every decision problem $\mathcal{D}$, tie-break function τ, every utility function u, and every factive and partitional evidence function $\mathcal{E}$,*

$$\sum_{E \in \mathcal{E}} C(E) \max_{o \in \mathcal{D}} \sum_{w \in W} s(l(C, E), \bar{o})(w)u(w) \geq \max_{o \in \mathcal{D}} \sum_{w \in W} s(C, \bar{o})(w)u(w)$$

where we abuse notation and write $\mathcal{E}$ to be the set of propositions you might learn—i.e. $\{E : \text{there is } w \text{ such that } E = \mathcal{E}_w\}$.

(ii) (a) *$l(C, E)(w) = C(w \mid E)$ whenever $C(E) > 0$; and*

(b) *$s(C, A)(w) = \sum_{w' \in W} C(w')f(A, w')(w)$, for some imaging operator f.*

It should be noted that Nielsen is here using a different account of the expected pragmatic value of an evidence-gathering episode from the one I have been using (Huttegger, 2013, 2014; Huttegger & Nielsen, 2020; Nielsen, 2024). Nielsen does not assume that we update by conditionalizing on the evidence we receive. Rather, he simply assumes that we have some way of updating. That is captured by the learning operator in his definition. Combined with an evidence function, this gives a function from worlds to posteriors: the evidence function gives a function from worlds to propositions, and then the learning operator, together with the prior, gives a function from propositions to posteriors. Nielsen then says that, from the point of view of a prior credence function, a particular learning operator combined with a particular evidence function satisfies the Value of Information property if the prior's expectation of the posteriors' expectation of the option the posterior would pick is at least as great as the prior's expectation of the option that the prior would pick.

This is different from the account I've been using so far, which says that, from the point of view of a prior credence function, a particular learning operator combined with a particular evidence function satisfies the Value of Information property if the prior's expectation of option the posterior would pick is at least as great as the prior's expectation of the option that the prior would pick. In the context of Savage's decision theory, we can show that, if the evidence function is factive and partitional, and we plan to update by conditionalization, then the two definitions are equivalent. And indeed we can also show that they are equivalent in the suppositional framework that Nielsen considers, providing we assume something about the suppositional expected value of a plan.

Here, I'm taking a plan to be a function that takes a world and returns the act you'll pick at that world. The plan that interests us here is the plan to pick whatever action maximizes suppositional expected utility relative to the posterior you'll have at that world once you receive the evidence and update. Let's call that plan *Learn*. What's its suppositional expected utility? Here's the assumption we need:

$$s(C, Learn) = \sum_{w} C(w) s(l(C, \mathcal{E}_w), \tau(\mathcal{D}_{l(C,\mathcal{E}_w)})$$

Risk-sensitive and ambiguity-sensitive decision theories

Shortly after Savage formulated the Sure Thing Principle, two challenges to it were raised: the so-called Ellsberg paradox and the so-called Allais paradox (Allais, 1953; Ellsberg, 1961). In each case we have a pair of decision problems, and when faced with those problems, people tend to choose in ways that are incompatible with the Sure Thing Principle. And yet we are hesitant to say that those choices are irrational. So there seem to be rational preferences that violate the Sure Thing Principle. And because they violate that, the Value of Information Theorem does not hold for them.

The Ellsberg paradox (Ellsberg, 1961). Before you sits an urn filled with red, black, and yellow balls. There are 90 balls in total. Exactly 30 are red. The remaining 60 are black or yellow, but you do not know in what proportion. A ball is about to be drawn from the urn. You are offered two different choices: the first between Option A and Option B; the second between Option C and Option D. The outcomes of these options depend on the colour of the ball. They're given in the following table:

	Yellow	Black	Red
Option A	£0m	£0m	£1m
Option B	£0m	£1m	£0m
Option C	£1m	£0m	£1m
Option D	£1m	£1m	£0m

Experimentally, Ellsberg found that people strictly prefer A to B, and D to C. The most common explanation is that they are averse to ambiguity. People prefer gambles for which they know the objective expected values. So they prefer A to B because they know the objective probability of Red (one-third) and the objective probability of Yellow or Black (two-thirds); and that is enough to calculate the objective expected value of A, but it does not fix the objective expected value of B, since the objective probability of Black could be anywhere from zero to two-thirds. Similarly, they prefer D to C because they know the objective probability of Red (one-third) and the objective probability of Yellow or Black (two-thirds); and that is enough to calculate the objective expected value of D, but not to calculate the objective expected value of C, since the objective probability of Yellow or Red could be anywhere from one-third to one, while the objective probability of Black could be anything from zero to two-thirds.

These preferences—often referred to as the Ellsberg preferences—violate the Sure Thing Principle. To see this, we reason by reductio. So suppose you have the Ellsberg preferences and you do satisfy the Sure Thing Principle. If you learn that the ball drawn is yellow, you'll be indifferent between A and B since they have the same outcome in that situation. So that means that, if you learn the ball drawn is not yellow, the Sure Thing Principle demands you must strictly prefer A to B, since you prefer A to B unconditionally. But now look at the second pair of options. If you learn that the ball drawn is Yellow, you'll be indifferent between C and D since they have the same outcome in that situation. So that means that, if you learn the ball drawn is not yellow, the Sure Thing Principle demands you must strictly prefer D to C, since you prefer D to C overall. But, if you learn the ball isn't yellow, then A is guaranteed to have the same outcome as C and B is guaranteed to have the same outcome as D. And so, if you learn the ball isn't Yellow, you should prefer A to B iff you prefer C to D. But that isn't true. We have a contradiction. And so you don't satisfy the Sure Thing Principle.

And we can translate this violation of the Sure Thing Principle into a failure of the Value of Information. Suppose (i) you prefer A to B, (ii) if you learn the ball is yellow, you'll be indifferent between A and B, and (iii) if you learn the ball isn't yellow, you'll strictly prefer B to A. Then you'll pay to avoid learning whether or not the ball is yellow. And similarly, suppose (i) you prefer D to C, (ii) if you learn the ball is yellow, you'll be indifferent between C and D, and (iii) if you learn the ball isn't yellow, you'll strictly prefer C to D. Then again you'll pay to avoid learning whether or not the ball is yellow. But, as we saw above, one of these must be true.

One upshot of this is that decision theories designed to accommodate the Ellsberg preferences will not always value gathering evidence, even when the evidence function is factive and partitional. One such decision theory is Γ-Maximin (Berger, 1985). Suppose your evidence rules out some objective probability functions, but not all. Then Γ-Maximin says you should pick an option whose minimal possible objective expected value among the objective probability functions compatible with your evidence is maximal.

The Allais paradox (Allais, 1953). You hold a ticket for a lottery in which there are 100 tickets. Before you learn which number is on our ticket, you are offered two different choices: the first between Option A′ and and Option B′; the second between Option C′ and Option D′. The outcomes of these options depend on the number on your ticket. They're given in the following table:

	Ticket 1-89	Ticket 90	Ticket 91-100
Option A′	£1m	£1m	£1m
Option B′	£1m	£0m	£5m
Option C′	£0m	£1m	£1m
Option D′	£0m	£0m	£5m

Allais claimed that many people prefer A′ to B′ and D′ to C′, and moreover that this pair of preferences is rational. This idea is that, faced with the first choice, a risk-averse person will prefer A′ over B′ because, while B′ gives a possibility of even greater wealth, the possibility of getting nothing makes B′ less desirable than A′ overall—A′ at least gains you something for sure. On the other hand, when faced with the choice between C′ and D′, where there is no option that guarantees a gain, the greater gain made possible by D′ seems worth it.

Again, and by exactly analogous reasoning, we can see that these preferences—the so-called Allais preferences—violate the Sure Thing Principle. And from that we can infer that, in either the choice between A′ and B′ or in the choice between C′ and D′, they do not value gathering the evidence whether the ticket number is between 1 and 89 or between 90 and 100.

One upshot of this is that decision theories designed to accommodate the Allais preferences will not always value gathering evidence, even when the evidence function is factive and partitional. These include John Quiggin's (1982; 1993) *rank-dependent utility theory*, Lara Buchak's (2013) *risk-weighted expected utility theory*, and Chris Bottomley and Timothy Luke Williamson's (ta) *weighted-linear utility theory*.

2.6 Generalizing the pragmatic Value of Information theorem

Now, the Value of Information Theorem as I have introduced it is severely limited in application: (1) evidence is rarely free; (2) inquiry involves not only deciding whether or not to gather a specific sort of evidence, but whether to gather this piece of evidence or that piece or to do something entirely different; (3) we rarely know exactly which decision we will face using our credences after the evidence is gathered; (4) evidence doesn't always tell you which member of a pre-specified partition is true; and (5) we'd like some reassurance that, when we do learn whatever we learn, the Bayesian command to update by conditioning on the evidence is the right one. In this section, we address these shortcomings.

#1: Factoring in the cost of evidence. While the Value of Information Theorem is interesting, the real value of the framework that Hosiasson, Blackwell, and Good introduced lies in the account of the pragmatic utility of an evidence-gathering episode at a world. And so, if there is a cost to gathering a certain sort of evidence, we can simply subtract the utility of whatever it is that it will cost us from the utility that gathering it gains for us to give the true pragmatic utility of gathering a specific piece of evidence; this even works if the cost is different in different worlds. And then we can take the expectation of this true pragmatic utility that factors in the cost, and compare it to the pragmatic utility of not gathering the evidence, which we can usually assume is cost-free.

#2: Comparing different evidence-gathering episodes. This account of the true pragmatic utility of gathering some evidence allows us to compare the expected utility of gathering *that* evidence with *that* cost to the expected utility of gathering *this alternative* evidence with *this* cost. After all, in inquiry, our choices are rarely simply to gather some evidence or not; they are choices between different evidence we might gather as well as other sorts of options; and the different sorts of evidence we might gather might have different costs. So, for instance, I might go to the window to see how the sky looks to inform my decision whether or not to take an umbrella, or I might look at the weather app on my smartphone, or I might do both, and each of these options might have different attendant costs.

In fact, Hosiasson's original paper, and Good's later one, both address a version of this question. Suppose you might gather evidence in one way, and it will teach you what the temperature is in Baltimore; or you might gather evidence in another way, and it will teach you something strictly more informative, such as what the temperature is in Baltimore *as well as* what the temperature is in Boston. Then, if the two evidence-gathering episodes have the same cost, the expected utility of engaging in the more informative one is at least the expected utility of engaging in the less informative one; what's more, that expected utility will be greater if you assign a positive credence to a world in which acquiring the less informative evidence would lead you to choose an option that is suboptimal from the point of view of the credences you would obtain at that world if you were to gather the evidence at that world.

As well as cases in which we wish to compare different evidence-gathering episodes we might undertake, we might also wish to compare an evidence-gathering episode and an alternative option that isn't an evidence-gathering episode at all. And the account of pragmatic value of evidence-gathering that we obtain from Hosiasson allows us to compare them as well. Perhaps I could check the sky from the window, check the weather app, do both, or I could do something else completely, such as making a sandwich for lunch. I can compare the expected utility of each and choose on that basis.

This allows us to consider the so-called opportunity cost of gathering a particular piece of evidence. This is not a cost that we factor into the pragmatic utility of each evidence-gathering episode. Rather, the opportunity cost incurred by doing one thing is the utility we would have got if we'd done some other thing instead. So the opportunity cost of gathering some evidence when I could have made a sandwich for lunch is whatever utility I would have got from making that sandwich. And the opportunity cost of gathering this evidence rather than that is the utility I would have got if I'd gathered that evidence instead.

#3: Allowing uncertainty about the decision problem you'll face. To define the pragmatic utility of an evidence-gathering episode, Hosiasson assumes you know for sure which decision you'll face using your credences, but of course you might be uncertain about this. Fortunately, it's easy to incorporate this uncertainty and recover a version of the Value of Information Theorem. We begin by ensuring that our possible worlds specify not only the truth values of the propositions to which we assign credences, but also which decision we'll face with our credences. We then assign credences to these more fine-grained possible worlds. And, having done all this, we define the pragmatic value of a credence function at a fine-grained world to be the utility at that world of the option it would lead us to choose from the decision we face at that world *once it updates on what decision problem we face at that world*. This latter clause is crucial: if we omit it, and fail to update on the decision problem we face when we become clear about which it is, then the Value of Information Theorem fails. So, the pragmatic value of an evidence-gathering episode is the pragmatic value of the credence function you'll end up with after gathering the evidence, updating on it, and learning what decision problem you face, and updating on that. With this amendment, the Value of Information Theorem still goes through.

Uncertainty about the decision problem

- Suppose $\mathbf{D}$ is a finite set of possible decision problems you might face.
- Suppose your credence function C is defined over the full algebra of subsets of $W \times \mathbf{D} = \{(w, \mathcal{D}) : w \in W \;\&\; \mathcal{D} \in \mathbf{D}\}$.
- Then define

$$\begin{aligned} \text{Exp}_C(\text{PU}^\tau(\mathcal{E})) &= \sum_{\substack{w \in W \\ \mathcal{D} \in \mathbf{D}}} C(w \;\&\; \mathcal{D}) \tau(\mathcal{D}_{C(-|\mathcal{E}_w \;\&\; \mathcal{D})})(w) \\ &= \sum_{\mathcal{D} \in \mathbf{D}} C(\mathcal{D}) \sum_{w \in W} C(w \mid \mathcal{D}) \tau(\mathcal{D}_{C(-|\mathcal{E}_w \;\&\; \mathcal{D})})(w) \end{aligned}$$

and

$$\begin{aligned}\text{Exp}_C(\text{PU}^\tau(C)) &= \sum_{\substack{w \in W \\ \mathcal{D} \in \mathbf{D}}} C(w \mathbin{\&} \mathcal{D})\tau(\mathcal{D}_{C(-|\mathcal{D})})(w) \\ &= \sum_{\mathcal{D} \in \mathbf{D}} C(\mathcal{D}) \sum_{w \in W} C(w \mid \mathcal{D})\tau(\mathcal{D}_{C(-|\mathcal{D})})(w)\end{aligned}$$

Then

Theorem 4. *If $\mathcal{E}$ is factive and partitional, and C is a probabilistic credence function, then*

(i)

$$\text{Exp}_C(\text{PU}^\tau(\mathcal{E})) \geq \text{Exp}_C(\text{PU}^\tau(C))$$

(ii)

$$\text{Exp}_C(\text{PU}^\tau(\mathcal{E})) > \text{Exp}_C(\text{PU}^\tau(C))$$

if there is w in W and $\mathcal{D}$ in $\mathbf{D}$ such that $C(w \mathbin{\&} \mathcal{D}) > 0$ and $\tau(\mathcal{D}_{C(-|\mathcal{D})}) \notin \mathcal{D}_{C(-|\mathcal{E}_w \mathbin{\&} \mathcal{D})}$.

The following example shows why we must insist that the individual update on the decision problem they face before choosing.

Example 2. *Suppose:*

- $W = \{w_1, w_2, w_3\}$;
- $\mathcal{E}_{w_1} = \{w_1, w_2\}$ $\quad \mathcal{E}_{w_2} = \{w_1, w_2\}$ $\quad \mathcal{E}_{w_3} = \{w_3\}$;
- $\mathcal{D} = \{o_1, o_2\}$ $\quad \mathcal{D}' = \{o'_1, o'_2\}$.

	w_1	w_2	w_3
$C(-)$	1/3	1/3	1/3
$C(- \mid \mathcal{E}_{w_1})$	1/2	1/2	0
$C(- \mid \mathcal{E}_{w_2})$	1/2	1/2	0
$C(- \mid \mathcal{E}_{w_3})$	0	0	1
$C(- \mid \mathcal{D})$	1	0	0
$C(- \mid \mathcal{D}')$	1/4	3/8	3/8
$o_1(-)$	−2	3	−2
$o_2(-)$	0	0	0
$o'_1(-)$	1	1	1
$o'_2(-)$	0	0	0

Then:

- $C(- \mid \mathcal{E}_{w_1})$ *will pick* o_1 *when faced with* $\mathcal{D}$
- C *will pick* o_2 *when faced with* $\mathcal{D}$*;*

- *C is certain of w_1 conditional on $\mathcal{D}$;*
- *so, C strictly prefers not learning to learning conditional on $\mathcal{D}$.*

What's more:

- *C, $C(- \mid \mathcal{E}_{w_1})$, $C(- \mid \mathcal{E}_{w_2})$, and $C(- \mid \mathcal{E}_{w_3})$ will all pick o'_1 when faced with $\mathcal{D}'$;*
- *so, C is indifferent between learning and not learning conditional on $\mathcal{D}'$.*

Therefore, since C assigns some positive credence to $\mathcal{D}$,

- *C strictly prefers not learning.*

#4: Allowing non-factive, non-partitional evidence. As stated, the Value of Information Theorem only covers cases in which the evidence-gathering episode will teach you which element of a partition is true. This is very idealized, though it is faithful to a certain way in which we gather evidence in science. When I measure the weight of a chemical sample, or when I ask how many organisms in a given population are infected after exposure to a particular pathogen, there is a fixed partition from which my evidence will come: I'll learn the sample is this weight or that weight or another one; I'll learn the number of infected organisms was zero or one or two or...up to the size of the population. But of course there are many cases in which our evidence-gathering will not be partitional or even factive in this way. Does the Value of Information theorem hold for any evidence function? If not, can we find weaker conditions on evidence-gathering episodes such that the theorem still holds?

First, we can easily see that there are evidence functions for which the Value of Information theorem does not hold. Take, for instance, a case in which you'll learn the opposite of what is true. If it's going to rain, you'll learn it won't, and if it won't, you'll learn it will. Then it's not hard to see that lots of priors will consider it worse in expectation to gather this evidence than not to, and for lots of decision problems. But there are less extreme cases, and in particular there are ones in which the evidence you receive is factive, unlike in the case just given. Here's one: You know the handkerchief in my pocket is rose, teal, or ochre, but you've no further information, so you assign the same credence to each. You can ask me what colour it is, and if it's rose, I'll say it's rose or teal, if it's teal, I'll say it's rose or teal, and if it's ochre, I'll say it's teal or ochre. And, later you'll face the following decision: you can choose a gamble, which will gain you three units of utility if it's teal and lose two if it's rose or ochre; or you can choose the sure thing, which will win you nothing and lose you nothing whatever colour it is. Then your current credences expect themselves to choose better in this situation than they expect your future credences to choose should you gather the information and update on it. To see this, note that your current credences will take the sure thing, whereas whatever you learn from the evidence-gathering, your updated credences will take the gamble.

Example 3 (The Colour of the Handkerchief)**.** *Suppose:*

- $W = \{w_1 = Rose, w_2 = Teal, w_3 = Ochre\}$;
- $\mathcal{E}_{w_1} = \{w_1, w_2\} = Rose \vee Teal$;
- $\mathcal{E}_{w_2} = \{w_1, w_2\} = Rose \vee Teal$;
- $\mathcal{E}_{w_3} = \{w_2, w_3\} = Teal \vee Ochre$;
- $\mathcal{D} = \{o_1, o_2\}$.

	w_1	w_2	w_3
$C(-)$	1/3	1/3	1/3
$C(- \mid \mathcal{E}_{w_1})$	1/2	1/2	0
$C(- \mid \mathcal{E}_{w_2})$	1/2	1/2	0
$C(- \mid \mathcal{E}_{w_3})$	0	1/2	1/2
$o_1(-)$	−2	3	−2
$o_2(-)$	0	0	0

Then, C will pick o_2, while $C(- \mid \mathcal{E}_{w_1})$, $C(- \mid \mathcal{E}_{w_2})$, and $C(- \mid \mathcal{E}_{w_3})$ will all pick o_1. So, C prefers not to learn the evidence.

So the Value of Information Theorem does not hold for all possible evidence-gathering scenarios, and even fails for reasonably quotidian ones—we'll see more examples below. So the question arises whether we can weaken the conditions of factivity and partitionality to understand better when it does hold. I'll describe three such attempts:

John Geanakoplos (1989) gives conditions on the evidence-gathering episode itself, and shows that, if it satisfies those, then for any prior credence function you have and any decision problem you'll face, providing you plan to update your prior by conditioning on whatever evidence you learn, gathering the evidence is never worse and often better than not gathering, in expectation. Nilanjan Das (2023) does something similar.

Geanakoplos' strengthening of the Value of Information theorem

Some preliminary definitions: Suppose $\mathcal{E}$ is an evidence function. Then:

- $\mathcal{E}$ is *factive* if $w \in \mathcal{E}_w$, for all w in W.

  That is, whatever evidence you receive will be true.

The evidence in Example B below isn't factive.

- $\mathcal{E}$ is *positively introspectible* if, whenever $w_2 \in \mathcal{E}_{w_1}$ and $w_3 \in \mathcal{E}_{w_2}$, then $w_3 \in \mathcal{E}_{w_1}$.

  That is, if your evidence at one world leaves another world open, and your evidence at the second world leaves a third world open, then your evidence at the first world should leave the third world open. But we can paraphrase it more intuitively as follows: if a certain proposition gives the strongest proposition you learn, then you also learn that the strongest proposition you've learned entails this proposition. After all, if $\mathcal{E}$ is positively introspectible, then, whatever world you're at, your evidence rules out all worlds at which the evidence you'd have there doesn't entail the evidence you actually have. So, if your actual evidence is E, then at all worlds at which E is true, your evidence at those worlds entails E.

  The evidence in Example C below isn't positively introspectible.

- $\mathcal{E}$ is *nested* if for any w_1 and w_2, either (i) $\mathcal{E}_{w_1} \subseteq \mathcal{E}_{w_2}$, (ii) $\mathcal{E}_{w_2} \subseteq \mathcal{E}_{w_1}$, or (iii) $\mathcal{E}_{w_1} \cap \mathcal{E}_{w_2} = \emptyset$.

  That is, if your evidence at two worlds overlaps, then one must entail the other.

  The evidence in Example C below isn't nested; nor is the evidence in the example of the Colour of the Handkerchief.

Example A below is factive, positively introspectible, and nested, but it is not partitional.

Theorem 5 ((Geanakoplos, 1989))**.** *If $\mathcal{E}$ is factive, positively introspectible, and nested, then for any prior credence function C, decision problem $\mathcal{D}$, and tie-breaking function τ,*

$$\text{Exp}_C(\text{PU}^{\mathcal{D},\tau}(\mathcal{E})) \geq \text{Exp}_C(\text{PU}^{\mathcal{D},\tau}(C))$$

with a strict inequality if there is w in W such that $C(w) > 0$ and $\tau(\mathcal{D}_C) \notin \mathcal{D}_{C(-|\mathcal{E}_w)}$.

Recall the example from above of the coloured handkerchief—I formalized it in Example 3. A key feature of that example is that the prior credences don't satisfy the Weak Reflection Principle with respect to the possible posteriors you might adopt after gathering the evidence.[2] This says that your prior credence function should be a weighted average of those possible evidence-

[2]The Weak Reflection Principle is my name for it. The principle itself is due to Bas van Fraassen (1999), who also formulated the Strong Reflection Principle we'll meet below.

informed posterior credence functions: that is, there should be some set of weights, one for each possible posterior, such that your credence in a given proposition is the weighted sum of the credences in that proposition assigned by the possible posteriors. You can tell that the credences in the handkerchief example violate this principle because the posterior credence in the handkerchief being teal is 1/2 whatever you learn in the evidence-gathering episode, while the prior credence is 1/3—no weighted sum of multiple 1/2s gives 1/3! The crucial fact is this: if your prior isn't a weighted sum of your possible posteriors, then there is a choice between two options in which your prior will choose one way, while all of your possible posteriors will choose the other way.[3] And in this case, it's clear that your prior will prefer not to gather evidence and stick with its own judgments rather than gather the evidence and thereby take on one of the posteriors to which that leads.

Weak Reflection Principle and the pragmatic value of information

A preliminary definition:

- An *updating plan* is a function R that takes a possible world w and returns a credence function R_w. The idea is that R_w is the credence function that R endorses at world w.

  If C is your prior, $\mathcal{E}$ is your evidence function, and you plan to update by conditionalization, then your updating plan will be $R_w(-) = C(- \mid \mathcal{E}_w)$. But clearly there are other possible updating plans.

- The pragmatic utility of an updating plan R at a world w is

$$\mathrm{PU}^{\mathcal{D},\tau}(R, w) = \mathrm{PU}^{\mathcal{D},\tau}(R_w, w).$$

Weak Reflection Principle Suppose C is your prior credence function and R is your updating plan. Then there should be weights $\lambda_w \geq 0$, one for each w in W, such that $\sum_w \lambda_w = 1$ and

$$C(-) = \sum_w \lambda_w R_w(-)$$

Theorem 6. *If C, R do not satisfy the Weak Reflection Principle, then there is a decision problem $\mathcal{D}$ such that, for any tie-breaking function τ,*

$$\mathrm{Exp}_C(\mathrm{PU}^{\mathcal{D},\tau}(R)) < \mathrm{Exp}_C(\mathrm{PU}^{\mathcal{D},\tau}(C))$$

In particular, $\mathcal{D} = \{o_1, o_2\}$ and

(i) *C prefers o_1 to o_2,*

(ii) *R_w prefers o_2 to o_1, for all w in W.*

[3] I state this result as Theorem 1 in (Pettigrew, 2023a), but it is not original to me.

So, if you violate the Weak Reflection Principle, there is a decision problem you might face such that your prior will prefer to make the decision itself rather than gather evidence and hand over to the resulting posteriors. But notice: satisfying the Weak Reflection Principle is certainly not sufficient to avoid this. Think of someone who is 50-50 about whether it will rain or not; if they choose to gather the evidence and it's raining, they'll learn it's not, but if they choose to gather it and it's not, they'll learn it is—so they're evidence is perfectly anti-factive. Their priors nonetheless satisfy the Weak Reflection Principle. So we need something stronger if we want a condition sufficient to make gathering evidence preferable to not gathering it.

One stronger condition that is sufficient is van Fraassen's original Reflection Principle, which I'll call the Strong Reflection Principle here (van Fraassen, 1984, 1995). This says that your prior credence function, conditional on your posterior being a particular probability function, is just that probability function. If your prior satisfies that principle with respect to your possible posteriors, then you will prefer to gather the evidence that gives your posteriors.

Strong Reflection Principle and pragmatic value of information

A preliminary definition:

- If R is your updating plan and P is in $\{R_w : w \in W\}$, then write $\overline{P}$ for the proposition that is true at all worlds w in W such that $P = R_w$. That is, $\overline{P}$ says that P is the credence function your updating plan endorses at your world.

Strong Reflection Principle Suppose C is your prior credence function and R is your updating plan. Then for all P in $\{R_w : w \in W\}$, the following should hold:

$$C(- \mid \overline{P}) = P(-)$$

if $C(\overline{P}) > 0$.

Theorem 7. *If C, R satisfy the Strong Reflection Principle, then, for any decision problem $\mathcal{D}$ and tie-breaking function τ,*

$$\text{Exp}_C(\text{PU}^{\mathcal{D},\tau}(R)) \geq \text{Exp}_C(\text{PU}^{\mathcal{D},\tau}(C))$$

with a strict inequality if there is w in W such that $C(w) > 0$ and $\tau(\mathcal{D}_C) \notin \mathcal{D}_{R_w}$.

However, as Kevin Dorst et al. (2021) note, there are credence functions that satisfy the Weak Reflection Principle but not the Strong Reflection Principle and yet for which the Value of Information Theorem holds. And we know there are credence functions that satisfy the Weak Reflection Principle for which the theorem does not hold. And so they seek a principle that lies somewhere between the two: stronger than the Weak version and weaker than the Strong one.

Enter the principle of Total Trust. Suppose we have what statisticians call a *random variable*. This is a quantity that can take different values at different possible worlds—it might be the amount of rainfall in Santiago next May, for instance. Now, suppose we take some threshold value—let's say 5mm—and consider the proposition that says that the expectation of this quantity from the point of view of your posterior credence function takes a value no less than this threshold—that is, it says that your posterior credence function expects the rainfall in Santiago next May to be at least 5mm. Now condition your prior on that proposition and, from the point of view of the resulting probability function, calculate the expected value of that quantity—so now we're talking about the expectation of the rainfall in Santiago next May from the point of view of your prior credence function once it's been conditioned on the proposition that your posterior expects that rainfall to be at least 5mm. Then Total Trust says that this expectation should also take a value no less than the threshold—your prior credence function conditional on the claim that your posterior will expect the rainfall to be no less than 5mm should expect the rainfall to be no less than 5mm. And this should hold for any quantity whatsoever—not just Santiago's rainfall. If your prior satisfies Total Trust, then the Value of Information Theorem holds.

Total Trust Principle and pragmatic value of information

Some preliminary definitions:

- Given an updating plan R, a random variable X, and a real number t, let $\langle \mathrm{Exp}_R(X) \geq t \rangle$ be the proposition that is true at all worlds w for which $\mathrm{Exp}_{R_w}(X) = \sum_{w' \in W} R_w(w')X(w') \geq t$.

Total Trust Principle Suppose C is your prior credence function and R is your updating plan. Then, for any random variable X and any threshold t, the following should hold:

$$\mathrm{Exp}_C(X \mid \mathrm{Exp}_R(X) \geq t) \geq t.$$

Theorem 8 ((Dorst et al., 2021))**.** *The following are equivalent:*

(i) *C, R satisfy the Total Trust Principle*

(ii) *For any decision problem $\mathcal{D}$ and tie-breaking function τ,*

$$\mathrm{Exp}_C(\mathrm{PU}^{\mathcal{D},\tau}(R)) \geq \mathrm{Exp}_C(\mathrm{PU}^{\mathcal{D},\tau}(C))$$

with a strict inequality if there is w in W such that $C(w) > 0$ and $\tau(\mathcal{D}_C) \notin \mathcal{D}_{R_w}$.

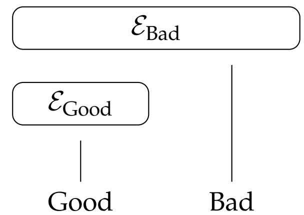

Figure 2.1: The evidence function in good/bad cases (Example A).

I won't delve any further into the details of the general results here; instead, I'll note a couple of examples that illustrate how many ways factivity and partitionality can fail while retaining the value of evidence-gathering.

A. Good and bad cases. In discussions of scepticism, whether it concerns the external world, other minds, or something else, externalists often distinguish themselves from internalists by claiming that the evidence you'd have in the 'good' or non-sceptical situation is different from the evidence you'd have in the 'bad' or sceptical situation (Williamson, 2013). In the good situation, your evidence is that you're in the good situation, while in the bad situation, your evidence is that you're either in the good situation or in the bad situation. Suppose I can gather evidence of this sort, perhaps by meeting another person about whom I am currently uncertain whether they have a mind, and then make a decision afterwards. Should I? The evidence in this case is factive, but non-partitional, since the evidence in the bad situation overlaps with the evidence in the good situation. Nonetheless, it satisfies the weaker conditions that Geanakoplos (1989) enumerates, and so it is always better in expectation to gather this evidence if you think it might lead you to change your mind about what to choose. Figure 2.1 illustrates the evidence function in this case.

B. Misdirection vs complete information. Someone in a company has committed fraud and it's your job to find out who it is. There are three suspects: the CEO, the COO, and the CFO. You have the opportunity to interview the CEO's assistant, and you know they know who did it. But you also know they're deeply loyal to the CEO. So, if it's the CFO, they'll tell you that; if it's the COO, they'll tell you that; but if it's the CEO, they'll tell you it's the CFO or the COO. So in this case, the evidence is not factive and it's not partitional. Figure 2.2 illustrates the evidence function in this case.

This is the sort of case that Nilanjan Das (2023) and Bernhard Salow (2018) call a *biased inquiry*, since there is a proposition such that you know your credence in it will rise regardless of what you learn, namely, the proposition that it was the CFO or the COO. This is a particular way in which you might violate the Weak Reflection Principle from above. And so we know that, for any prior that gives positive credence to each of the possibilities—CEO, COO, CFO—there is a decision problem your priors will prefer to face using themselves rather than the posteriors they'd get from gathering the evidence.

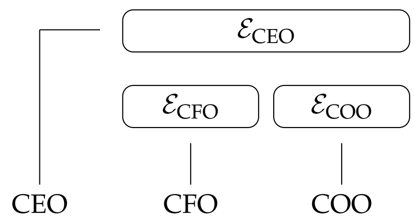

Figure 2.2: The evidence function in the CEO/COO/CFO case (Example B).

But that's not necessarily true of all decision problems. If you have a reasonably high prior that it's the CEO or if there's a big difference in the utilities of the different options at the state of the world at which it is the CEO, then you should not take the evidence, since it's too misleading relative to your prior and the decision problem. But if there's no difference between the utility of the options at the world at which the CEO is guilty, perhaps because you know there's no way to prosecute that individual anyway, then you should take the evidence, since it gives the opportunity of learning exactly who did it if it's the one of the other two. So this is a very clear case in which you have to weigh up misleading evidence, which you'll receive if the CEO is guilty, against highly accurate and informative evidence, which you'll receive if the CEO is innocent. How you weigh it up depends on your priors, but also the decision you face.

Example 4 (Fraud in the company)**.**

- $W = \{w_1 = CEO, w_2 = CFO, w_3 = COO\}$;
- $\mathcal{E}_{w_1} = \{w_2, w_3\} = CFO \vee COO$;
- $\mathcal{E}_{w_2} = \{w_2\} = CFO$;
- $\mathcal{E}_{w_3} = \{w_3\} = COO$.

	w_1	w_2	w_3
$C(-)$	1/3	1/3	1/3
o_1	1	4	2
o_2	1	2	3
o'_1	−3	1	1
o'_2	0	0	0

Then:

- *C prefers to learn the evidence before facing $\mathcal{D} = \{o_1, o_2\}$.*

  After all, there is no difference between o_1 and o_2 in the world in which they'll get false evidence, so it doesn't matter what they'll pick; and the evidence will lead them to pick the best for sure in the other worlds.

- *C prefers not to learn the evidence before facing $\mathcal{D}' = \{o'_1, o'_2\}$.*

  After all, C prefers o'_2 to o'_1, while each possible posterior prefers o'_1 to o'_2.

At this point, it might occur to you to ask: is there no limit to the falseness of evidence we might sometimes prefer to acquire? In the case just described, you weighed up the possibility of false evidence against the possibility of very informative true evidence. But could there be a case in which all the possible evidence is false and yet you'd still choose to gather it? The answer is yes. Suppose there are four states of the world, and you must choose between two options. The first gives zero units of utility for sure, while the second gives positive utility at two worlds and negative utility at two worlds. Then your ideal situation would be to have credences that choose the first option at the worlds at which the second has negative utility and the second option at worlds at which the second has positive utility. Now suppose that, if you're in one of the worlds at which the second option as negative utility, you'll learn you're at the other world at which it has negative utility; and if you're in one of the worlds at which the second option has positive utility, you'll learn you're at the other world at which it has positive utility. Then gathering the evidence before choosing will lead you to choose whichever option is best at whatever world you're in. And that's better, in expectation, than picking whichever of the two options maximizes expected utility from the point of view of your prior.

Example 5 (Anti-factivity)**.**

- $W = \{w_1, w_2, w_3, w_4\}$;
- $\mathcal{E}_{w_1} = \{w_2\}$ $\quad \mathcal{E}_{w_2} = \{w_1\}$ $\quad \mathcal{E}_{w_3} = \{w_4\}$ $\quad \mathcal{E}_{w_4} = \{w_3\}$.

	w_1	w_2	w_3	w_4
$C(-)$	1/4	1/4	1/4	1/4
o_1	1	2	−1	−1
o_2	0	0	0	0

Then:

- *C will pick o_1;*
- *If you gather the evidence, you'll pick o_1 in worlds w_1 and w_2, and o_2 in worlds w_3 and w_4, and doing that is exactly as good at worlds w_1 and w_2 as picking o_1, which C will do there, and it is better than picking o_1 in w_3 and w_4, which is again what C will do there, and so C prefers to gather the evidence.*

We can represent gathering evidence, updating on it, and picking in line with the updated credences as a new option on the list:

	w_1	w_2	w_3	w_4
$C(-)$	1/4	1/4	1/4	1/4
o_1	1	2	−1	−1
o_2	0	0	0	0
GATHER	1	2	0	0

And, since GATHER *weakly dominates o_1 and o_2, C prefers it.*

Reflecting on this example gives us an interesting way to understand why learning evidence can be better, in expectation, than not learning it. Essentially, the availability of evidence makes available a new option in the decision problem that isn't there if you don't gather the evidence. In the example just given, the available options were zero-utility-for-sure or positive-utility-at-two-worlds-negative-utility-at-two-worlds. But, the evidence described there made available a different option: positive-utility-at-two-worlds-zero-utility-at-two-worlds. It made it available because by choosing to gather the evidence and then decide, and knowing how you'll update and then choose, you are essentially choosing the option whose utility at a world is the utility of whatever option you'll choose if you first update on the evidence at that world and then choose using those credences. And in the case just described, the option is at least as good as each of the original options at every world and better than each at some. So it is better in expectation.

And, reflecting on this fact, we see that gathering evidence will have no value if the range of available options is rich enough. Suppose a coin has been tossed twice, giving four possible outcomes, HH, HT, TH, TT. And suppose you have the opportunity to learn whether the first coin landed heads—i.e. the disjunction HH or HT—or whether the first coin landed tails—i.e. the disjunction TH or TT. And suppose there are four possible options, which I'll write as the quadruple of their payouts at the four different worlds: (u_1, u_2, u_3, u_4), (u_1, u_2, v_3, v_4), (v_1, v_2, u_3, u_4), and (v_1, v_2, v_3, v_4). Then whatever my credences and whatever the values of these utilities, I will not pay to gain the evidence. For suppose my prior prefers (u_1, u_2, u_3, u_4) to the rest, and my posterior upon learning the first toss landed heads also prefers (u_1, u_2, u_3, u_4), while my posterior upon learning the first toss landed tails prefers (v_1, v_2, v_3, v_4). Then, if I were to prefer learning the evidence to avoiding it, I must prefer (u_1, u_2, v_3, v_4) to (u_1, u_2, u_3, u_4), since the former gives the pragmatic utilities of learning the evidence. But I don't have that preference, and I know I don't because that first option (u_1, u_2, v_3, v_4) is available to me, and I don't prefer it. And similarly for the other possibilities.

C. Williamson's unmarked clock. Externalists often contend that our evidence is not luminous to us; that is, we can have evidence that does not rule

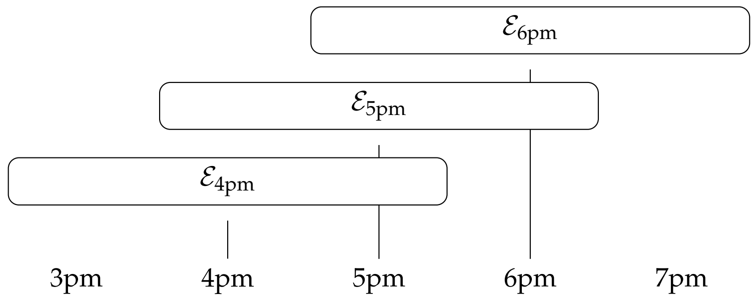

Figure 2.3: Partial illustration of the evidence function in the unmarked clock case (Example C).

out our evidence being different from how it actually is. A neat illustration is Tim Williamson's (2014) example of the unmarked clock. You want to know the time. You can walk through to the next room and look at a clock. But alas, it is a fashionable clock of the minimalist sort favoured at the moment, and it has no numbers marked on it. It has only a sweeping single hand. Before you enter the room, you know the hand is pointing at 1 or 2 or 3 or . . . or 10 or 11 or 12, but you know don't know which. If it's 1, your evidence will be that it's 12, 1, or 2; if it's 2, your evidence will be that it's 1, 2, or 3; and so on. That's how well your eyesight can discern the differences. Figure 2.3 illustrates part of the evidence function in this case.

Should you look at the clock? As in the example of corporate fraud above (Example B), it very much depends on your priors and the decision you'll face. Suppose you currently assign equal credence to each possible position of the sweeping hand. The first option available pays a million dollars if it's at 1 or 2 or . . . or 6, and nothing otherwise; the second pays a million dollars if it's at 7 or 8 or . . . or 12, and nothing otherwise. Then you should gather the evidence—you'll do better in expectation if you do. But, as Nilanjan Das (2023) notes, if the first option pays a million dollars if the hand points to an even number and the second pays a million dollars if it points to an odd number, then you shouldn't gather the evidence, since the evidence is misleading about whether the time is even or odd: if it's even, you'll become twice as confident it's odd as you are that it's even, and vice versa.

It is worth pausing here to consider a philosophical use to which Simon Huttegger (2013, 2014) wishes to put the value of information framework. Picking up a suggestion by Brian Skyrms (1990a), Huttegger wishes to suggest that the value of information property can distinguish genuine learning events from events that change our credences but not via learning. So, suppose you know that after some event your credences will be given by either this credence function or that credence function, but you know neither which it will be nor what the mechanism will be by which you'll come to have them—Skyrms and Huttegger describe the credal change event as a black-box. Then, they say, the event is a genuine learning experience just in case, for every decision problem you might face, you'll prefer to have the experience first and then choose using the credences you'll end up with rather than choose using your current credences. So the reason that the change in your credences that result from, say, taking a hallucinogenic drug does not count as a genuine learning experience is that, from the point of view of your current credences, you'd expected yourself to choose more poorly after the change than you expect yourself to choose now.

It seems to me that this gives an elegant criterion by which to distinguish certain sorts of credal change that we don't want to describe as genuine learning events from those we do—events like taking a hallucinogenic drug or getting hit on the head, for instance. But I worry that it sets too high a standard for a genuine learning experience. For instance, in the case of the unmarked clock, it seems that we do have a genuine learning experience: my evidence will be factive, and it will rule out quite a lot of options that I currently rule in, and so it's pretty informative. But nonetheless it fails Skyrms' and Huttegger's test.

#5: Assessing updating plans. Throughout, we have assumed that, whatever evidence we gather, we update on it in the Bayesian's standard way by conditioning on the proposition learned. But, as Peter M. Brown (1976) showed, we can use the Hosiasson-Blackwell-Good framework to argue for this Bayesian assumption, at least in those cases to which the Value of Information Theorem originally applied, namely, cases of factive and partitional evidence.

An *updating plan* is a function that takes a possible world and returns a credence function. You might think of the credence function as the one the plan endorses at that world. We can easily define the pragmatic utility, at a world, of an updating plan relative to a decision problem and tie-breaker function, to be the pragmatic utility, at that world, of the credence function it endorses at that world. Of course, what we'd most love is to follow the plan that takes each world to its omniscient credence function, that is, the one that assigns maximal credence to all truths and minimal credence to all falsehoods. But following that plan isn't available to us. Rather, we must pick a plan that assigns to two worlds the same credence function whenever those two worlds give rise to the same evidence. We'll call these the *available updating plans* relative to an evidence-gathering episode. Now, given an evidence-gathering episode, we can then ask which of the available updating plans has the greater expected pragmatic utility from the point of view of a prior credence function. Brown shows that, if the episode is factive and partitional, then updating plans that require you to condition on whatever evidence you learn maximize expected pragmatic utility.

What about cases in which the evidence is not factive or not partitional? Then Miriam Schoenfield (2017) shows that the updating plans that maximize expected pragmatic utility are not those that require you to condition on your evidence, but those that require you to condition on the fact you received the evidence you did. That is, in the unmarked clock case described above, if the sweeping hand points to 2 and I receive the evidence that it points to 1 or 2 or 3, I should conditionalize not on this evidence, but on the fact I received it, which is true only at the world at which it points at 2.

There is an interesting ongoing debate whether such an updating plan is really available to me (Carr, 2021; Gallow, 2021; Isaacs & Russell, 2023; Schultheis, ta). You might think it is not, since I could only implement it by reflecting on the evidence I in fact have, and inferring the worlds at which I would receive that evidence. But of course, in the cases at hand, I'm supposed not to know what evidence I have, and so presumably I can't reflect on it. However, externalists do think I should update by conditionalizing on the evidence I have, and you might think that this equally requires me to know what evidence I have.[4]

I won't delve deeper into this debate here, but I will note briefly that Gallow (2021) and Salow (2018) both argue that it should not be possible to enter yourself into evidential situations in which updating in the rationally required way is guaranteed to increase your credence in a false proposition. And yet that is exactly what happens in the unmarked clock case with respect to the proposition that the clock is pointing at an even number. This leads Salow to say that such evidential situations can't arise, offering an internalist account of evidence instead, and it leads Gallow to offer an alternative updating rule. My own take on this situation is that the evidential situation can arise and that you should update by conditionalizing when you face it, but the possibility of self-delusion to which this gives rise is not problematic because putting yourself in such a self-deluding evidential situation is rationally impermissible from the point of view of the value of information: you simply shouldn't view the clock if it is the accuracy of your credence in the odd/even proposition that you care about; and you shouldn't do it if the decision you'll face with your credences is a bet on that same proposition.

To wrap up this section, it is worth noting that, if Schoenfield's updating rule is genuinely available to us, then the Value of Information Theorem holds for *any* evidence function, whether factive, partition, only one, only the other, or neither. That is, if we assume that we'll respond to evidence not by conditioning on the evidence we learn but on the fact that we learn it, then gathering evidence is always at least as good in expectation as not gathering it, and it is strictly better in expectation if you think learning it might lead you to choose a different option when you face the decision problem.

Brown's and Schoenfield's pragmatic arguments for updating

Some preliminary definitions:

- Given an evidence function $\mathcal{E}$, an updating plan R is *available in* $\mathcal{E}$ if, whenever $\mathcal{E}_w = \mathcal{E}_{w'}$, $R_w = R_{w'}$.

  That is, your updating plan can't discern the world any more precisely than your evidence can; it cannot respond differently at

[4] Though Gallow (2021) describes an alternative updating rule the externalist might endorse instead.

worlds at which you receive the same evidence.

- Given an evidence function $\mathcal{E}$ and a world w, let $\overline{\mathcal{E}_w}$ be the proposition that is true at all worlds at which your evidence is the same as it is at world w.
- Given an evidence function $\mathcal{E}$ and a prior C, an updating plan R is a Schoenfield plan for C and $\mathcal{E}$ if $R_w(-) = C(- \mid \overline{\mathcal{E}_w})$, whenever $C(\overline{\mathcal{E}_w}) > 0$.

Theorem 9 ((Brown, 1976; Schoenfield, 2017))**.** *Suppose $\mathcal{E}$ is an evidence function, C is a prior credence function and R, R' are updating plans. Then*

(i) *If R is a Schoenfield plan for C and $\mathcal{E}$, and R' is an available plan in $\mathcal{E}$, then, for any decision problem $\mathcal{D}$ and tie-breaker function τ,*

$$\text{Exp}_C(\text{PU}^{\mathcal{D},\tau}(R)) \geq \text{Exp}_C(\text{PU}^{\mathcal{D},\tau}(R'))$$

(ii) *If R is a Schoenfield plan for C and $\mathcal{E}$, and R' is an available plan in $\mathcal{E}$ that is not a Schoenfield plan for C and $\mathcal{E}$, there is a decision problem $\mathcal{D}$ and tie-breaker function τ such that*

$$\text{Exp}_C(\text{PU}^{\mathcal{D},\tau}(R)) > \text{Exp}_C(\text{PU}^{\mathcal{D},\tau}(R')).$$

Notice that, if $\mathcal{E}$ is factive and partitional, and R is a Schoenfield plan for C and $\mathcal{E}$, then $R_w(-) = C(- \mid \mathcal{E}_w)$, since $\mathcal{E}_w = \overline{\mathcal{E}_w}$.

Also notice that since the trivial updating plan on C, which takes every world and returns C, is always available, and it corresponds to not gathering the evidence at all, this theorem shows that, regardless of how your evidence function is, if you will update on new evidence using a Schoenfield rule, then gathering evidence is always at least as good in expectation as not, and it is strictly better, in expectation, if you think learning the evidence might lead you to change your mind about how to choose.

Chapter 3

The epistemic value of information

3.1 The epistemic value of gathering evidence

So far, we have focused on the pragmatic version of the Value of Information theorem, which tells us something about when you have *practical* reason to engage in a certain sort of evidence-gathering. But, as Graham Oddie (1997) showed, and Wayne Myrvold (2012) generalized, there is also a version that tells us something about when you have *epistemic* reason to gather evidence. Alejandro Pérez Carballo (2018) has extended Oddie's and Myrvold's approach in various ways, and we will see that Dorst et al. (2021) provide insights analogous to those they provided in the pragmatic case.

Recall: Hosiasson's insight is that the pragmatic value of a credence function is the utility of the option it leads you to choose, and the pragmatic value of an episode of evidence-gathering is the pragmatic value of the credence function it will lead you to have after you update your prior on the evidence you learn. But credence functions don't just have pragmatic value. We don't use them only to guide our decisions. We also use them to represent the world, and their purely epistemic value derives from how well they do that, regardless of whether we need them to help us choose.

Many ways of measuring this purely epistemic value have been proposed, but by far the most popular characterizations of the legitimate epistemic utility functions say that they are all *strictly proper*, where this means that, if we measure epistemic utility in this way, any probabilistic credence function expects itself to have strictly greater epistemic utility than it expects any alternative credence function to have; that is, it thinks of itself as uniquely best from the epistemic point of view; that is, it is epistemically immodest. Jim Joyce (2009) defends something close to this view, and Robbie Williams and I (2023) have recently argued for it in a different way. What's more, it is widely assumed throughout accuracy-first epistemology, which seeks to understand and ground the epistemic normativity of credences by investigating the optimal ways in which to pursue epistemic value understood as credal accuracy.[1]

[1]For an overview of this approach to epistemic normativity, see (Pettigrew, 2023c).

Strictly proper epistemic utility functions

An *epistemic utility function* EU takes a credence function C and a possible world w and returns $\text{EU}(C, w)$, a real number or ∞ or $-\infty$, which measures the epistemic value of C at w.

Definition 3 (Strict propriety)**.** EU *is strictly proper if, for any probabilistic credence function P and any alternative credence function $C \neq P$,*

$$\text{Exp}_P(\text{EU}(P)) = \sum_{w \in W} P(w)\text{EU}(P, w) > \sum_{w \in W} P(w)\text{EU}(C, w) = \text{Exp}_P(\text{EU}(C))$$

Perhaps the most well-known strictly proper epistemic utility function is the so-called *Brier score*. Given a proposition, we say that the omniscient credence in it is 1 if it's true and 0 if it's false. The Brier score of a credence function at a world is then obtained by taking each proposition to which it assigns a credence, taking the difference between the credence it assigns to that proposition and the omniscient credence in that proposition at that world, squaring that difference, taking the average of these squared differences, and then subtracting the result from 1.

In the Brier score, each proposition is given equal weight in the average, but we can also give greater weight to some propositions than others in order to record that we consider them more important. This gives a *weighted Brier score*. This is important in the current context, since it allows us to explain why it is better, epistemically speaking, to engage in some evidence-gathering episodes rather than others, even when the latter will improve certain credences more than the former will improve others. The explanation is that the credences the latter will improve are less important to us. So, for instance, one evidence-gathering episode might, in expectation, greatly improve the accuracy of my credences concerning how many blades of grass there are on my neighbour's lawn, while another might, in expectation, only slightly improve the accuracy of my credences about the fundamental nature of reality, and yet I might favour the latter because the propositions it concerns are more important to me.

Another strictly proper epistemic utility function, less well-known but interesting nonetheless, is the *enhanced log score*. If a proposition is true, we score a credence in it by subtracting that credence from its own logarithm; if a proposition is false, we score a credence in it by subtracting that credence from zero. The enhanced log score of a credence function is then the average of these scores across all credences it assigns, and a weighted enhanced log score is a weighted average of them.

The Brier score and the enhanced log score

Definition 4 (Brier score)**.** *The Brier score* Brier(C, w) *of a credence function* C *at* w *is*

$$\text{Brier}(C, w) = 1 - \frac{1}{n} \sum_{X \in \mathcal{F}} |C(X) - V_w(X)|^2$$

where $V_w(X) = 1$ *if* X *is true at* w *and* $V_w(X) = 0$ *if* X *is false at* w*, and* n *is the number of propositions in* $\mathcal{F}$*.*

To give a weighted Brier score, we assign to each proposition X *in* $\mathcal{F}$ *a weight* $0 < \lambda_X < 1$*, where* $\sum_{X \in \mathcal{F}} \lambda_X = 1$*, and then define it as follows:*

$$\text{Brier}_\Lambda(C, w) = 1 - \sum_{X \in \mathcal{F}} \lambda_X |C(X) - V_w(X)|^2$$

The Brier score and any weighted Brier score are strictly proper.
In the diagram below, we plot the Brier score of a single credence in a true proposition in red (i.e. $y = 1 - (1-x)^2$), and the Brier score of a single credence in a false proposition in blue (i.e. $y = 1 - x^2$).

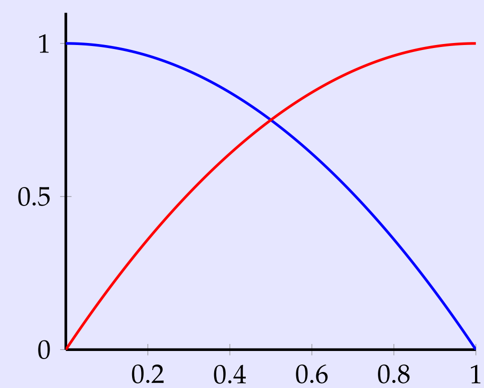

Definition 5 (Enhanced log score)**.** *The enhanced log score* Log(C, w) *of a credence function* C *at* w *is*

$$\text{Log}(C, w) = \frac{1}{n} \sum_{X \in \mathcal{F}} V_w(X) \log(C(X)) - C(X)$$

where again $V_w(X) = 1$ *if* X *is true at* w *and* $V_w(X) = 0$ *if* X *is false at* w*, and* n *is the number of propositions in* $\mathcal{F}$*.*

To give a weighted enhanced log score, we assign to each proposition X *in* $\mathcal{F}$ *a weight* $0 < \lambda_X < 1$*, where* $\sum_{X \in \mathcal{F}} \lambda_X = 1$*, and then define it as follows:*

$$\text{Log}_\Lambda(C, w) = \sum_{X \in \mathcal{F}} \lambda_X [V_w(X) \log(C(X)) - C(X)]$$

The enhanced log score and any weighted enhanced log score are strictly proper.

In the diagram below, we plot the enhanced log score of a single credence in a true proposition in red (i.e. $y = \log(x) - x$), and the enhanced log score of a single credence in a false proposition in blue (i.e. $y = -x$).

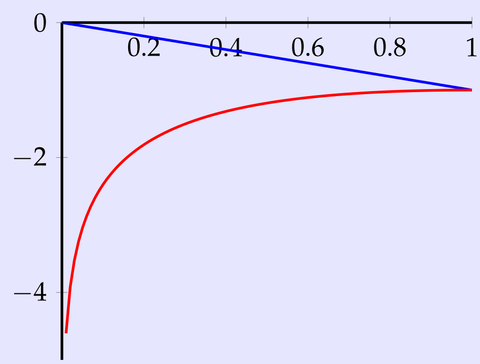

So now we have a way of assigning epistemic value to a credence function at a world. And so we can say that the epistemic value, at a world, of gathering evidence is the epistemic value of the credence function you'll end up with when you update on the evidence you'll get at that world—as before, we begin by assuming you update by conditioning on your evidence. And now we can state Oddie's epistemic version of the Value of Information Theorem: suppose you may gather evidence that will teach you which element of a particular partition is true, and suppose your epistemic utility function is strictly proper; then the expected epistemic value of gathering the evidence, from the point of view of your current credences, is always at least as great as the expected epistemic value of not gathering the evidence, from the same point of view; and, if you give some positive credence to a state of the world at which what you will learn will lead you to change your credences, then the expected epistemic value of gathering the evidence is strictly greater than the expected epistemic value of not doing so.

Oddie's Epistemic Value of Information Theorem

Theorem 10 ((Oddie, 1997))**.** *If* EU *is strictly proper and* $\mathcal{E}$ *is factive and* *partitional,*

$$\mathrm{Exp}_C(\mathrm{EU}(\mathcal{E})) = \sum_{w \in W} C(w)\mathrm{EU}(C(- \mid \mathcal{E}_w), w) \geq \sum_{w \in W} C(w)\mathrm{EU}(C, w) = \mathrm{Exp}_C(\mathrm{EU}(C))$$

with strict inequality if there are w, w' *such that* $\mathcal{E}_w \neq \mathcal{E}_{w'}$ *and* $C(w), C(w') > 0$.

One thing that often surprises people about this result is that it seems to contradict the definition of strict propriety. According to strict propriety, every probabilistic credence function thinks it's best; but now we learn that it thinks that gathering evidence and updating on it to give different credence functions is even better. What's going on?

In fact, there is no contradiction: each probabilistic credence function thinks that it is better, in expectation, than any other specific credence function; but the updating plan isn't a specific credence function—it's different credence functions at different worlds. And strict propriety doesn't rule out a probabilistic credence function preferring a strategy that gives different credence functions at different worlds. Take, for example, the strategy, unavailable to all but God, of simply adopting, at a world, the omniscient credence function at that world, that is, the credence function that gives maximal credence (i.e. credence 1) to propositions that are true at that world and minimal credence (i.e. credence 0) to those that are false. Then this strategy gives the best credence function at each world. And so any credence function thinks of this strategy as better than itself, in expectation. But that doesn't contradict strict propriety.

As with the pragmatic version of the Value of Information Theorem, the reason Oddie's result holds is not that learning true evidence is guaranteed to improve your epistemic situation, and so certainly will improve it in expectation. As before, it's quite possible to acquire true evidence that is misleading. For instance, suppose my credence it's sunny is 10%, my credence it's windy is 40%, and my credence it's rainy is 50%. And suppose it's sunny. I then learn it's sunny or windy and my credence in sun becomes 20% and my credence in wind becomes 80%. Then, according to the Brier score, my epistemic utility dropped from 0.59333 to 0.57333. So my evidence was misleading and my epistemic situation deteriorated as a result of learning true evidence. But, as in the pragmatic case, Oddie's result holds because, in the particular conditions he places on the evidence-gathering episode, it will always be the case that any epistemic deterioration, once weighted by the prior's probability that it will happen, is outweighed by the epistemic improvements that are possible, once those are weighted by the prior's probability that they will happen instead.

3.2 Generalizing the epistemic Value of Information theorem

As with the Value of Information Theorem, we can generalize this result. As long as we set up an exchange rate between epistemic and pragmatic utility, we can factor in the cost of the evidence. That is, once we say how much pragmatic utility we're prepared to pay for a given amount of epistemic utility, we can say when gathering evidence is the right thing to do, rationally speaking. And, as before, we can use the expected epistemic utilities of different evidence-gathering episodes, with their costs factored in, to choose between them, and choose between them and doing something entirely different, which doesn't involve gathering evidence at all. And finally, we can generalize beyond factive and partitional evidence in a similar way.

Geanakoplos-style strengthening of Oddie's Theorem

Theorem 11 ((Dorst, 2020; Dorst et al., 2021; Levinstein, 2023))**.** *If $\mathcal{E}$ is factive, positively introspectible, and nested, then for any prior credence function C and any strictly proper epistemic utility function* EU*,*

$$\mathrm{Exp}_C(\mathrm{EU}(\mathcal{E})) \geq \mathrm{Exp}_C(\mathrm{PU}(C))$$

with strict inequality if there are w, w' such that $\mathcal{E}_w \neq \mathcal{E}_{w'}$ and $C(w), C(w') > 0$.

Simple Trust Principle and the epistemic value of information

Some preliminary definitions:

- An epistemic utility function EU is *additive and continuous* if there is a function $s : \{0,1\} \times [0,1] \to [-\infty, \infty]$ such that $s(1, x)$ and $s(0, x)$ are continuous functions of x, and

$$\mathrm{EU}(C, w) = \sum_{X \in \mathcal{F}} s(V_w(X), C(X)).$$

- Given an updating plan R, a proposition X, and a real number $0 \leq t \leq 1$, let $\langle R(X) \geq t \rangle$ be the proposition that is true at all worlds w for which $R_w(X) \geq t$.

Simple Trust Principle Suppose C is your prior credence function and R is your updating plan R. Then, for any proposition X and any threshold t, the following should hold:

$$C(X \mid R(X) \geq t) \geq t.$$

Theorem 12 ((Levinstein, 2023))**.** *The following are equivalent:*

(i) *C, R satisfy the Simple Trust Principle*

(ii) *For any additive and continuous strictly proper epistemic utility function* EU,

$$\mathrm{Exp}_C(\mathrm{EU}(R)) \geq \mathrm{Exp}_C(\mathrm{EU}(C))$$

with strict inequality if there is w such that $R_w \neq C$ and $C(w) > 0$.

What about the cases we considered above?

A. Good and bad cases. In such a case, relative to any strictly proper scoring rule, this evidence will increase your epistemic utility in expectation. This is no surprise: in the bad world, your credence function will remain the same after learning the evidence; in the good world, it will become perfectly omniscient; and so in expectation, learning will be an improvement.

B. Misdirection vs complete information. In this case, relative to the Brier score, there are priors that will expect this information to increase epistemic utility and priors that will expect it to decrease. If your prior credence that the CEO did it is high enough, then there is no credences for CFO and COO that would making learning desirable; if your prior credence that the CEO did it is low (below $1/2$, say), then learning can be desirable if your credences in CFO and COO are unequal enough.

Interestingly, relative to the enhanced log score, learning the evidence is never desirable. The reason is that, at the world at which it's the CEO, your credence function will assign credence zero to the true possibility, and this has epistemic utility $-\infty$ (since the logarithm of zero is negative infinity), and so the expected epistemic value of gathering the evidence is $-\infty$, whereas for any credence function its expected epistemic utility by its own lights is always greater than $-\infty$.

C. Williamson's unmarked clock. Again, relative to the Brier score, there are priors that will expect this information to increase epistemic utility and priors that will expect it to decrease. For instance, if you have equal credence in each of 1 or 2 or ... or 11 or 12, then you increase your Brier score in expectation by gathering the information, while if you lump nearly all of your credence onto one of the numbers, you decrease it. And this time, the same is true for the enhanced log score.

3.3 Assessing updating plans

As in the pragmatic case, we can appeal to measures of epistemic value to assess updating plans; and, when we do, we get the same results that we got in the pragmatic case. If your evidence function is factive and partitional, and your epistemic utility function is strictly proper, you maximize expected epistemic utility by choosing to update by conditionalization (Greaves & Wallace, 2006). And, more generally, regardless of what your evidence function is like, the available updating plan that maximizes expected utility relative to any strictly proper epistemic utility function is the one that tells you to condition not on your evidence but on the fact that you learned that evidence; that is, you maximize expected epistemic utility by choosing to update by Schoenfield conditionalization (Schoenfield, 2017).

Greaves & Wallace's and Schoenfield's epistemic arguments for updating

Theorem 13 ((Greaves & Wallace, 2006; Schoenfield, 2017)). *Suppose $\mathcal{E}$ is an evidence function, C is a prior credence function and R, R' are updating plans. Then:*

(i) *If R is a Schoenfield plan for C and $\mathcal{E}$, and R' is an available plan in $\mathcal{E}$, then, for any strictly proper epistemic utility function* EU,

$$\text{Exp}_C(\text{EU}(R)) \geq \text{Exp}_C(\text{EU}(R'))$$

(ii) *If R is a Schoenfield plan for C and $\mathcal{E}$, and R' is an available plan in $\mathcal{E}$ that is not a Schoenfield plan for C and $\mathcal{E}$, then, for any strictly proper epistemic utility function* EU,

$$\text{Exp}_C(\text{EU}(R)) > \text{Exp}_C(\text{EU}(R')).$$

Recall: If $\mathcal{E}$ is factive and partitional, then the Schoenfield plans are precisely the conditionalization plans.

3.4 Combining the pragmatic and epistemic values of information

I conclude this tour of the arguments and results concerning the value of information by noting that, once you set your exchange rate between the epistemic and the pragmatic, you can specify the all-things-considered value of an evidence-gathering episode. So, for instance, suppose your epistemic utility function is the Brier score; and suppose you fix a particular scale on which you'll measure your pragmatic utility (for recall that pragmatic utility is equally well measured on any of an infinite collection of scales, each obtained from any other by a positive linear transformation). Then, in order to specify the all-things-considered value of an evidence-gathered episode, you need to know how much pragmatic utility you consider equal to, say, 0.1 change in Brier score. If it's 0.1 units of pragmatic utility (measured on the scale we fixed), then your all-things-considered utility is just the sum of your pragmatic utility (on the fixed scale) and the Brier score. If it's 0.2, then your all-things-considered utility is your pragmatic utility added to double the Brier score—0.1 change in Brier score is worth twice a 0.1 change in pragmatic utility. And so on.

This all-things-considered utility allows us to incorporate both the practical value of having credences as well as their representational value. Both are important to us, and the trade-off may well be important in certain cases. There will be cases in which learning the evidence increases our Brier score in expectation, but decreases the pragmatic utility of our credences because of the decision problem we will face with them. For instance, in the case of Williamson's unmarked clock, if we have equal prior credences in 1 or 2 or . . . or 11 or 12, then we increase our Brier score by looking at the clock, but decrease the pragmatic utility of our credences for sure if the bet we're going to face is whether the hand points at an odd or even number. In that case, then, we must determine our exchange rate to discover whether looking at the clock is all-things-considered the right thing to do.

Part II

The epistemology of inquiry

Chapter 4

When should we gather evidence?

In the second part of this essay, I turn from the formal results and arguments that extend Hosiasson's and Good's Value of Information Theorem and focus on the literature on the epistemology of inquiry that has been developing recently in mainstream epistemology. My plan is to apply the insights from the first part to see what light they might shed on some of the central questions that have arisen in that literature.

The ambition of this part is imperialistic or, in Eric Schliesser's (ta) happier terminology, totalizing. In Schliesser's conception of synthetic philosophy, we take a totalizing approach when we take a formal framework from one area and apply it across a very wide range of cases in that area and in other areas: it might be the framework of natural selection, for instance, or the Bayesian treatment of uncertain belief. In my case, the framework is our standard theory of rational decision-making under uncertainty, developed in philosophy, economics, and statistics during the twentieth century.

Applied to the normativity of inquiry, this framework results in the value of information approach I described in the first part of this essay. Decisions whether or not to inquire are just a particular species of decision under uncertainty, and so they fall within the ambit of rational choice theory. To apply this theory to this sort of decision, we need an account of the value of a particular zetetic act, such as an evidence-gathering episode or a sequence of them. In the pragmatic case, this is given by Hosiasson's insight that the value of gathering evidence when you'll face a particular decision problem is just the value of the option you'll choose from those available in that decision problem with the credences you'll have after gathering the evidence. In the epistemic case, as Oddie teaches us, it's just the epistemic value of the credences you'll have after gathering the evidence. With these accounts in hand, we can apply rational choice theory to zetetic decisions just as we apply it to other decisions under uncertainty. And we can use it to say when it is pragmatically or epistemically or all-things-considered rational to start, continue, conclude, and reopen inquiry.

The totalizer claims, furthermore, that this approach says everything there is to say about the rationality or correctness or appropriateness of doing these things. Rational choice theory is, after all, supposed to be a complete theory of the rationality of such choices.

To see the approach in action, let's ask four of the central questions in the epistemology of inquiry: when should we embark on a particular inquiry? when should we continue to pursue an inquiry on which we've already embarked? when should we conclude one? when should we reopen an inquiry we previously concluded?

4.1 What are the acts of inquiry?

When we ask these questions, we immediately face a further question that always arises when we use rational choice theory: how extended are the different actions that are available to the decision-maker, in this case, the potential inquirer? Take the example of a detective on a murder case trying to decide whether or not to inquire. Can they choose now to undertake the full inquiry, an extended action that involves first viewing the crime scene, then running the forensic tests, then taking the witness and suspect statements, then searching the suspects' homes? This is an action that might take a couple of weeks to complete: is it possible for them, at one point in time, to choose to perform that whole act? Or are they choosing to undertake only the first of these evidence-gathering episodes, namely, viewing the crime scene, but in the knowledge that taking that first step and learning its outcome will affect whether they continue with that inquiry or whether they move to another, or to something else entirely? Or are they choosing an even less extended action, such as walking towards the crime scene, or taking the first step in that walk?

Decision theorists have various ways of determining which actions are actually available to a decision-maker at the time of the decision. You might feel that the least extended action—the half-second action of taking the first step—is clearly available; perhaps the more extended action of viewing the crime scene—which we might say will take fifteen minutes—is less obviously available, but nonetheless is sufficiently available; and the fully extended, fortnight-long action of carrying out all of the inquiry is not sufficiently available to the potential inquirer at the very first point of the inquiry. One way to make this intuition clear is to say that the degree of availability of an action is the probability that the decision-maker would fully enact it were they to choose to do so. If the detective were to decide to walk towards the crime scene, it's very likely they would undertake and complete that full act; if they were to decide to view the crime scene, it's a bit less likely, but still very likely they would undertake and complete it; but if they were to decide to undertake the whole action, there are so many things that might knock them off course or force them to reconsider before the action is complete, that it is not sufficiently likely they would undertake and complete it.

For our purposes in what follows, we will choose the grain at which we describe the available actions to suit the examples we're considering. When we're not interested in whether the individual might be knocked off course even in the execution of a single evidence-gathering episode, we'll simply take the whole episode to be an available action.

Whatever we choose, when we ask whether we should embark on an inquiry, we are asking whether we should undertake the first available action in that inquiry, which might be an evidence-gathering episode or it might be something shorter, such as an initial step in such an episode. And when we evaluate such an action, we must factor into the decision how likely it is that we'll proceed to the next step of the inquiry should we embark on the first, how much utility the outcome of that second step would have in conjunction with the outcome of the first, how likely it is we'll proceed to the next, as well as its utility, and indeed how likely it is that we'll complete the inquiry and what would be the utility of the outcome of doing that. So the options available to the potential inquirer are not *undertake-and-complete-the-inquiry* and *do-not-undertake-and-complete-the-inquiry*; they are usually *undertake-the-first-evidence-gathering-episode-of-the-inquiry* and *do-not-undertake-the-first-evidence-gathering-episode-of-the-inquiry*, where of course the second option might split further into many other possibilities, such as pursuing a different inquiry or doing something else entirely.

4.2 When should we begin, continue, or restart a particular inquiry?

So, how should you choose whether or not to begin an inquiry, or continue pursuing one you've already begun, or reopen one you previously closed? The totalizing answer is this: in the way you should choose everything else! From the pragmatic point of view, you should do so if doing so maximizes subjective expected pragmatic utility; from the epistemic point of view, you should do so if doing so maximizes subjective expected epistemic utility; and from the all-things-considered point of view, you should do so if doing so maximizes subjective expected all-things-considered utility.

An interesting consequence of this: since an inquiry is a series of evidence-gathering episodes, it can be rational to embark on it even if not all of the episodes that make it up lead to improvements in the expected pragmatic or epistemic value of your credences, and even if some of the episodes lead to a decrease in those expected values, just as it can be rational to embark on a series of dental procedures even though you know that some of the individual procedures in the series will make things worse. Provided you're confident enough that you'll see the series through to the end, and provided the series in full leads to sufficiently great improvements in expectation, and provided your dental situation wouldn't be too much worse if the series got interrupted in the middle, it is rational to embark on it.

Sometimes, we embark again on an inquiry we have already completed: we double-check our results. On the face of it, this seems a puzzling practice. After all, we've undertaken the inquiry, and we've concluded it to our satisfaction. Why, then, are we undertaking it again? Woodard (2022) addresses this challenge in the knowledge framework, where the worry takes this form: if you have concluded an inquiry into a question, then you have come to know the answer, but you should not embark on an inquiry into a question unless you are ignorant of the answer, and so you should not inquire twice into the same question. She maintains in the face of this objection that it is permissible to double-check, and illustrates her point with a series of vignettes in which people appear to double-check rationally.

The Value of Information framework allows us to appreciate a number of circumstances in which it is rationally permissible, and perhaps even required to double-check:

In one sort of case, what we call double-checking is better described as acquiring a second sample of evidence from the same source. This is what we do if we check the door is locked when we leave our building, but then go back to double-check a minute later (cf. Woodard's example of Deming). We are not carrying out the same evidence-gathering episode twice; rather, we are carrying out two different but very closely related evidence-gathering episodes: Does the door open when I pull the handle this time? Does the door open when I pull the handle this different time? The first episode will already make me very confident the door is locked. What's more, the outcomes of the two episodes are very highly correlated with one another: that is, the probability of the door not opening when I pull it a second time, given it didn't open when I pulled it the first time, is extremely high and much higher that the unconditional probability of the door not opening on the second try. And so the second evidence-gathering episode can't hope to shift my credence that it's locked up upwards by more than a tiny amount. However, if I take the stakes to be very high and the cost very low, it might nonetheless be rational.

In another sort of case, we double-check something for which we had very good evidence at some point in the past, but not because we want to get even better evidence, but because we want to regain the confidence given us by the very good evidence we previously had. This is what we do, for instance, if we packed our favourite condiment in our luggage three days ago, but then today we open the luggage again to double-check it's there (cf. Woodard's example of Sam). In this case, we had very good evidence that the hot sauce is in our luggage three days ago when we put it there. Our immediate perceptual experience provided that evidence. Perhaps it warranted a credence of 99%. But gradually, as this experience has moved from being immediate to being a memory, it has come to support only a lower credence by today. Perhaps now the memory of that experience warrants only a credence of 90%. Perhaps I'm wrong about what I seem to perceive immediately before me only one time in a hundred, but I'm wrong about what I seem to remember perceiving immediately before me one time in ten. And perhaps from the point of view of that lower credence, checking to see whether the hot sauce is there is something that maximizes expected utility, in which case double-checking is permitted, and maybe even required.

4.3 When should you cease inquiring further?

The totalizing view I am propounding here says that, just as you should embark on inquiry when doing so maximizes the subjective expectation of whatever variety of utility you're interested in, and similarly for continuing to pursue an inquiry and reopening a previously closed one, so you should cease inquiring further when continuing to inquire no longer maximizes that subjective expectation; that is, when there is an alternative available action that has greater subjective expected utility than continuing to inquire, whether that alternative available action is an evidence-gathering episode from a different inquiry, or some other action altogether, like making a sandwich.

Your reasons for gathering further evidence can just run out, and from that point of view it can be irrational to pursue your inquiry any further. For instance, this can happen if you care only about the pragmatic value of your credences as a guide to action in the face of a particular decision you know you'll face. At some point, you might come to know that all further evidence-gathering episodes that are actually available to you either won't change your mind about what to choose when faced with the decision, or that any that might change your mind are too costly. At this point, further inquiry is irrational from this myopic pragmatic point of view. While you might continue to improve your credences from an epistemic point of view, you achieve no further gains from a pragmatic point of view—at least none that issue from the decision you'll use those credences to face.

This can lead you to abandon before they're complete inquiries that it was nonetheless rational to embark on in the first place. For instance, this might happen because the costs of gathering further evidence in that inquiry have increased since you began the inquiry; or because the stakes of the decision you'll face using the credences you'll form have decreased; or because it becomes cheaper to inquire in a different way, a way that you thought would be too costly at the beginning of your inquiry, but which you have since learned is actually rather inexpensive—for instance, we can imagine a detective who is scouring through all CCTV footage over a 24 hour period because they thought that DNA testing would be expensive, but who has recently learned it's very cheap and so switches to that, abandoning their original inquiry.

From the epistemic point of view, things are a little different. Unless you somehow acquire certainty about the correct answer to the question at which your inquiry aims, there will always be some evidence-gathering episode that you'll expect to improve your credence function from a purely epistemic point of view, though of course that episode may not be available to you, or it might be too costly, and those are reasons to cease that inquiry.

Now, you will rarely acquire the sort of certainty that concludes inquiry regardless of what actions are available or what cost they have. After all, for most inquiries, the evidence-gathering episodes don't give definitive answers to the target question. They give definitive answers to related questions that bear on the target question, such as when I gather evidence about what the weather forecast says as part of my inquiry into whether or not it will rain tomorrow.

This vindicates a point raised by Avery Archer (2021) and Christopher Willard-Kyle (forthcoming), who argue that, in inquiry, there will nearly always be room for improvement from an epistemic point of view. They are responding to those who say that knowledge is the aim of inquiry, and that an inquiry concludes once the inquirer knows the answer to the defining question. They argue this can't be right because, even after you've achieved this knowledge, it's always possible to improve your epistemic situation. After all, you might obtain *better* knowledge of the correct answer: you might obtain a safer belief, even though your current belief is sufficiently safe to count as knowledge; or you might obtain the belief you currently have, but using an even more reliable process, even though your current belief was formed by a sufficiently reliable process; and so on.

One interesting possibility that this throws up is that even those who think it is knowledge and not mere accuracy that we value will need to provide something like a numerical measure of the epistemic value of a doxastic state; to wit, an epistemic utility function. After all, one upshot of Willard-Kyle's point is that someone who knows the answer to the question at which their inquiry is aimed must decide whether or not to continue to pursue this inquiry. As he points out, by doing so, they can continue to improve their epistemic situation, but presumably there are diminishing marginal returns from such efforts, and so they must weigh those expected gains against the expected gains brought by some other pursuit. So, for instance, the detective who now knows that the suspect was at the scene of the crime can continue to inquire about that in order to improve the quality of her knowledge of it, or she can turn her attention to another question, such as whether they have a motive. To choose between these two courses of action, she must be able to weigh the improvements that each will bring in expectation. And that requires some way of measuring their epistemic value. I won't pursue this any further.

The main claim of the totalizing view, then, is that you should treat zetetic decisions in the same way you treat other decisions: you may do them when they are among the options that maximize subjective expected utility, you should do them when only they do this, and you should refrain from doing them when they don't. They are, after all, simply choices to do certain things; and their outcomes can be evaluated for their utility in exactly the same way the outcomes of other choices can; and our uncertainty about which outcome will eventuate can be treated as it can for other decisions. It is true that there is a purely epistemic perspective from which we sometimes wish to assess an evidence-gathering episode, but we saw in Chapter 3 that we can accommodate that as well; and we can combine it with the pragmatic perspective to give the all-things-considered perspective. With all of this in hand, we can now turn to some of the views that have been developed in the recent literature on the epistemology of inquiry.

Chapter 5

From evidence-gathering to reasoning

So far, we have asked only when we should gather empirical evidence, which empirical evidence we should gather, what price we should pay for it, and when our reasons for gathering it run out. But inquiry comprises more than simply gathering empirical evidence. Once gathered, we can do various things with this evidence: we can reason from it alone, or we can combine it with other evidence we have and reason from that. And indeed, we can reason from no evidence at all, as we do in logic or in mathematics or in conceptual pursuits, such as philosophy. The reasoning we undertake is sometimes logical, such as when we bring together our evidence that it's either going to be windy or wet later, our evidence that if it's windy I'll need to bring in the washing from the line, and our evidence that if it's wet I'll need to bring in the washing from the line, and reason from those three pieces of evidence to the conclusion that I'll need to bring in the washing and form a high credence in that; the reasoning is sometimes conceptual, such as we'll see below when someone brings together their evidence that two streets are parallel with their evidence that one runs in a particular direction to infer that the other does too and to form a high credence in that; it is sometimes inductive, such as when we bring together each individual observation we've made about a particular phenomenon to draw an inductive generalization from them; it is sometimes abductive, such as when we consider a body of evidence, formulate a hypothesis, convince ourselves that it explains the evidence better than anything else available and infer to that; and so on. Since such reasoning is a central part of how we should proceed under uncertainty, the Bayesian's imperialistic or totalizing ambitions require them to produce a theory of such reasoning that allows us to represent it within the Bayesian framework and provide norms to govern it. And doing so, we will see, also allows the Bayesian to answer some of Jane Friedman's (2020) recent questions about when we should gather empirical evidence and when we should carry out reasoning from our existing evidence.

5.1 Epistemic and zetetic norms in conflict

Let me begin with Friedman's central example in order to illustrate the problem that our Bayesian theory of reasoning from evidence can solve. After that, I'll present such a Bayesian theory, based on an insight by Ian Hacking (1967) that Robbie Williams (2018) and I (2020b) have developed further elsewhere, and I'll explain what it says about Friedman's case.

Here is my version of Friedman's example:

> **D. Jane the Glazier.** Jane is a glazier who hopes to win the contract to replace all the windows in the Chrysler building in New York City. In order to formulate her bid, which would be extremely lucrative should it succeed, she needs to know how many windows there are in the building. So she hotfoots it down to Grand Central Station and stands there counting the windows. This takes time and it takes concentration. While Jane is deliberately and intentionally engaged in this evidence-gathering episode, which will issue in empirical evidence about the number of windows in the building, she inadvertently and unintentionally receives other evidence that is irrelevant to that particular inquiry. For instance, she overhears someone next to her say to their companion that Madison Avenue runs from north-east to south-west (henceforth, NE-SW). She receives this evidence, and responds as she does to much testimony by becoming certain of that proposition. However, she doesn't draw out all of its consequences; that is, she doesn't reason from it in all the ways she might; indeed, she doesn't reason from it in any of the ways she might. For instance, a few months ago, while idly looking at a map of Manhattan, Jane noticed that Madison Avenue and Park Avenue run parallel to one another. So she could reason from that evidence, together with the evidence she's just inadvertently received from the person next to her outside Grand Central Station, to the conclusion that Park Avenue runs NE-SW. But to do that would require time and effort; it would require concentration and attention. And she needs to keep her attention on counting the windows if she is to complete that inquiry, which is very important to her.

Friedman concludes that, in cases like this, there is a tension between epistemic norms and zetetic norms. Here are two of the epistemic norms she mentions:

> **EP**$_a$ If one has excellent evidence for p at t, then one is permitted to judge p at t.
>
> **EP**$_o$ If one has excellent evidence for p at t, then one ought to judge p at t.

And here is the zetetic norm:

> **ZIP** If one wants to figure out [the answer to a particular question], then one ought to take the necessary means to figuring out [that answer].

Jane the glazier has excellent evidence that Park Avenue runs NE-SW—after all, she has evidence Madison Avenue runs in that direction and she has evidence they're parallel. And yet, were she to do what is necessary to draw out that conclusion, she'd divert attention from counting the windows. So, ZIP clashes with EP_a and EP_o in this case: ZIP says she must continue to concentrate on counting; EP_o says she's required not to; EP_a says she's permitted not to.

What's more, it's clear that none of the norms just given are true in general. The case of Jane the glazier provides a counterexample to EP_a and EP_o. But we might easily change it to provide a counterexample of ZIP. Suppose, for instance, she overhears the person next to her describing some symptoms they've just started experiencing—numbness in their face, blurred vision in one eye, a headache, difficulty walking. She responds by raising her credence in the proposition that this person is experiencing these things, but she also intuitively feels a sense that these might be concerning without knowing exactly why, but knowing that, if she were to attend to what she's just learned, she would be able to figure out what's happening with this person and how to help them. In that version of the case, while she indeed wants to figure out how many windows there are in the Chrysler building, it is of greater import to reason out the consequences of her evidence about this person's medical condition, and so she ought not to take the necessary means to figuring out how many windows there are in the Chrysler building. So ZIP is false: she must reason in this case, and abandon the inquiry in which she is engaged.

So one task for our Bayesian account of reasoning from your evidence is to say why Jane should continue counting the windows in the original case and why she should stop counting them in the case in which she overhears the person describing the symptoms of a stroke. And indeed, more generally, our task is to replace EP_a, EP_o, and ZIP with a true general norm that governs these situations.

5.2 A Bayesian account of reasoning

Let me now turn to that Bayesian account of reasoning from your evidence. This is based on a suggestion that Ian Hacking (1967) made, and that Robbie Williams and I have developed in different ways (Williams, 2018; Pettigrew, 2020b).

Hacking developed this account in order to address the problem of old evidence. Suppose someone formulates a hypothesis about some phenomenon in the world—gravitation, for instance. On the standard Bayesian account, a piece of evidence *supports* or *confirms* that hypothesis if the conditional probability of the hypothesis given the evidence is greater than the unconditional probability of the hypothesis.

The standard Bayesian account of confirmation

$$E \text{ confirms } H \text{ iff } P(H \mid E) > P(H).$$

However, often we are already certain of a piece of evidence before we become aware of a hypothesis that it might support. And, in that case, the conditional probability of the hypothesis given the evidence is just identical to the unconditional probability of the hypothesis. And yet surely it is possible that what we might call 'old evidence'—evidence of which we are certain before we become aware of a hypothesis—can confirm that hypothesis. To cite a common example, certain observations that had been made before Einstein formulated the theory of relativity are nonetheless taken to confirm that theory. But it seems the Bayesian account of confirmation must deny this.

Enter Hacking. Take a case in which a hypothesis entails a piece of old evidence. He thought that what happens in these cases is that, to show that the evidence confirms the hypothesis, you consider the conditional probability of the hypothesis given not only the old evidence but also the logical fact that the hypothesis entails the old evidence. And if that conditional probability is greater than the unconditional probability of the hypothesis, then the evidence confirms the hypothesis.

Hacking's account of confirmation

Suppose H entails E (written $H \Rightarrow E$). Then:

$$E \text{ confirms } H \text{ iff } P(H \mid E \ \& \ (H \Rightarrow E)) > P(H).$$

However, the fact that the hypothesis entails the evidence is a logical fact. And, according to the orthodox classical version of Bayesianism, you are rationally required to be certain of all logical facts. And so, if you are certain of some piece of evidence, then the conditional probability of the hypothesis given not only that evidence but also the logical fact that the hypothesis entails it is identical to the unconditional probability of the hypothesis. So we've made no progress.

But, Hacking contends, the orthodox classical Bayesian account is wrong, and he offers an alternative on which you are not required to be certain of all logical truths. You can rationally be less than certain of them, and you can come to learn them, and when you come to learn them that can lead you to increase your credence in other propositions, and so it is possible to be certain of some evidence but less than certain that a hypothesis entails that evidence, and so come to raise your credence in the hypothesis upon learning that the hypothesis entails the evidence.

To present Hacking's account, let's begin with classical Bayesianism.

When I introduced the Bayesian account of uncertain thought in the first part of these notes, I adopted what I called a very flat-footed approach. I assumed that credences take *propositions* as their objects, and indeed I assumed that propositions are simply sets of possible worlds. To understand Hacking's theory, we need to take a more nuanced approach. So I will assume here that credences in fact take *declarative sentences* as their objects, and sentences are the sort of thing that can be true or false. This allows us to represent someone who assigns different credences to 'It's raining' and to 'It's raining or it's raining', even though these sentences are both true at exactly the same possible worlds, namely, the set of all worlds in which it's raining.

So let's focus on a particular individual who assigns credences to a particular set of sentences—we'll call that set of sentences their agenda. In the flat-footed framework in which an individual's credences are assigned to sets of possible worlds, we assumed that our individual assigns a credence to each possible world—or, more precisely, they assign a credence to each atomic proposition, which is a proposition true at exactly one possible world. And then we stated the standard Bayesian norm of Probabilism by saying that (i) their credences in all of these possible worlds (that is, in all the corresponding atomic propositions) should add up to 1, and (ii) their credence in a particular proposition should be the sum of their credences in the possible worlds (that is, their credences in the atomic propositions corresponding to those worlds) at which the proposition is true.

In the more nuanced framework we assume here, we also want to talk about possible worlds. But instead of taking the possible worlds as primary and taking the propositions to be sets of those possible worlds, we take the sentences in the individual's agenda to be primary in this framework and take the possible worlds to be classically consistent assignments of truth values to those sentences. We don't assume that our individual assigns credences to these possible worlds, or to sentences true at exactly one of these worlds. But we do appeal to something similar to state the classical version of Probabilism in this framework. Classical Sentential Probabilism says that there must be a set of non-negative numbers, one assigned to each possible world—that is, to each classically consistent assignment of truth values—such that (i) these numbers sum to 1, and (ii) our individual's credence in a sentence should be the sum of the numbers assigned to the possible worlds—that is, the classically consistent truth value assignments—on which that sentence is true. These numbers are called probability weights and we say that they generate the credences in question. The more standard characterization of classical probabilism follows from this: if a sentence is a classical tautology, then it's true at all classically consistent assignments of truth values, and so it should receive the sum of all the probability weights, which is 1; if a sentence is a classical contradiction, then it's false at all classically consistent assignments of truth values, and so it should receive the sum of none of the probability weights, which is 0; if two sentences are classically inconsistent and a third is true at all classical assignments at which one or other of the first two sentences is true, then the third should receive the sum of the weights of the assignments at which it's true which is the sum of the weights of the assignments at which the first is true plus the sum of the weights of the assignments at which the second is true.

And Classical Sentential Conditionalization says that, if our individual learns something that rules out certain possible worlds, the numbers that they previously assigned to the worlds they've ruled out should be set to zero, and the numbers assigned to the worlds they haven't ruled out should be increased by the same factor so that they now sum to 1, and the new credence our individual assigns to a sentence should be the sum of the new numbers assigned to those possible worlds at which it's true. Again, the more standard characterization of classical conditionalization follows: if there is a sentence to which our individual assigns a credence that is true at all and only the worlds not ruled out by whatever the individual learns, then Classical Sentential Conditionalization requires that the posterior credence in a sentence is the prior credence in the sentence conditional on the sentence learned.

Classical Sentential Bayesianism

Some notation:

- $\mathcal{L}$ is the set of sentences to which our individual assigns credences;
- $V_{\mathcal{L}}$ is the set of classically consistent assignments of truth values to the sentences in $\mathcal{L}$. If A is in $\mathcal{L}$ and w is in $V_{\mathcal{L}}$ and A is true at w, we write $w \models A$.
- $\Lambda = \{\lambda_w : w \in V_{\mathcal{L}}\}$ is a set of *probability weights* if (i) $0 \leq \lambda_w \leq 1$, for all w in $V_{\mathcal{L}}$, and (ii) $\sum_{w \in V_{\mathcal{L}}} \lambda_w = 1$.
- Given a set of probability weights Λ and a credence function C, we say that Λ *generates* C if, for any A in $\mathcal{L}$,
$$C(A) = \sum_{\substack{w \in V_{\mathcal{L}} \\ w \models A}} \lambda_w$$
- Given a set of probability weights Λ and a subset $E \subseteq V_{\mathcal{L}}$, we write Λ^E for the set of probability weights $\{\lambda^E_w : w \in V_{\mathcal{L}}\}$ where
$$\lambda^E_w = \begin{cases} 0 & \text{if } w \notin E \\ \frac{\lambda_w}{\sum_{w' \in E} \lambda_{w'}} & \text{if } w \in E \end{cases}$$

Classical Sentential Probabilism If C is our individual's credence function, there should be a set of probability weights Λ such that Λ generates C.

Classical Sentential Conditionalization Suppose:

(i) C is our individual's credence function at time t,

(ii) C' is her credence function at a later time t'.

Then there must be a set of probability weights Λ that generates C such that Λ^E generates C'.

This might seem rather a convoluted and unconventional route to stating the classical Bayesian norms in the case where the objects of credence are sentences. But it pays dividends because it allows us to weaken Classical Sentential Bayesianism to give Hacking's more permissive theory in a very straightforward way. I'll call Hacking's theory Personal Sentential Bayesianism. It looks identical to the classical version I've just presented except that the assignments of truth values are no longer required to be classical. That is, to obtain Hacking's version of Bayesianism, we vastly enlarge the set of what counts as a possible world so that it includes *all* assignments of truth values, whether they are classically consistent or not. So, for instance, there will be a possible world where 'It's raining' is true but 'It's raining or it's raining' is false; there will be one in which the sentence 'There are infinitely many prime numbers' is false; there will be one in which 'It's raining' and 'If it's raining, then it's wet' are true, but 'It's wet' is false; and so on.

Personal Sentential Bayesianism

Let $V^*_{\mathcal{L}}$ be the set of all assignments of truth values to the sentences in $\mathcal{L}$. And replace $V_{\mathcal{L}}$ by $V^*_{\mathcal{L}}$ throughout the formulation of Classical Sentential Bayesianism. The result is Personal Sentential Bayesianism.

One way to motivate Personal Sentential Bayesianism is to think about the arguments we give in favour of Probabilism. The betting argument seeks to establish that, if you have credences that don't obey Probabilism, then there is a set of bets those credences will lead you to accept that, taken together, are guaranteed to lose you money (Ramsey, 1926 [1931]; de Finetti, 1931; Vineberg, 2016; Pettigrew, 2020a). The accuracy argument seeks to establish that, if you have credences that don't obey Probabilism, then there is an alternative set of credences over the same agenda as yours that are guaranteed to be more accurate than yours (Joyce, 1998; Predd et al., 2009; Pettigrew, 2016, 2023c). How compelling this argument is depends on how we understand what it means for something to be guaranteed. If the betting and accuracy arguments are to establish Classical Probabilism, it must mean that the bets in question lose you money at every logically possible world and the credences in question are more accurate at every logically possible world. But, if that's what we mean, the argument isn't very compelling. What's so irrational about doing something that is guaranteed to lose you money at a set of possibilities that's more restricted than the set of possibilities that are possible for you at that time? After all, it might make you money at worlds outside this set. For instance, let's suppose that Goldbach's conjecture in number theory is true, and therefore logically necessary. But, since it is not yet proved, I am less than certain of it. That means, I'll pay a certain price for a bet against it. Now, this bet will lose me money at all logically possible worlds—but of course it will win me money at those logically impossible worlds at which Goldbach's conjecture is true. This doesn't render me irrational, for those logically impossible worlds are still open for me, and not because of any error or failure on my part, but just because I haven't proved the conjecture (and nor has anyone else). Similarly, there will be a credence function more accurate than mine at every logically possible world—but it will be less accurate at those logically impossible worlds at which Goldbach's conjecture is true. Again, that doesn't seem to render me irrational.

To drive the point home, suppose I've never heard that Cary Grant and Archibald Leach were the same person. So I'm less than certain of this. Then there is a bet I'll accept that will lose me money at all metaphysically possible worlds, and there are alternative credences I might have that are more accurate at all metaphysically possible worlds. Now, surely that doesn't render my credences irrational. But, if a *metaphysical* guarantee of lost money and less accuracy does not render me irrational, why should a *logical* guarantee of these things do so? Logical truths, like metaphysical truths, are things we have to discover. We might not discover them in the same way, and in some sense it might be true that we have the resources we need to discover the logical truths inside our head all along, and so don't need to investigate the world to discover them in the way we do with certain metaphysical necessities like the identity of Cary Grant and Archibald Leach. But that doesn't change the fact that we must do something—logical or mathematical reasoning—to discover them. And it doesn't change the fact that, as finite beings who can only do certain things and not everything, our failure to discover all logical truths is no failure on our part. Hacking's goal in expanding the set of possibilities to include all the assignments of truth values, not just the logically possible ones, is to ensure that, if you were to disobey the resulting version of Probabilism, the betting and accuracy arguments for that version would really have bite.

In Personal Sentential Bayesianism, it's permissible to be less than fully certain of logical facts, since it's permissible to have a credence function generated by probability weights that give some weight to possible worlds in which the logical fact is false. Such an assignment of truth values is of course not classically consistent, but that doesn't matter because in Hacking's framework classically inconsistent truth value assignments are among the possible worlds. What's more, we can now model logical reasoning in this framework as analogous to gathering empirical evidence. In both cases, what we do is rule out possibilities. When empirical evidence comes in the form of a sentence we learn—such as when Jane learns 'Madison Avenue runs NE-SW'—it rules out those possible worlds in which that evidence is false. When we do logical reasoning, that also rules out possible worlds. We see this most clearly when we conduct our logical reasoning by constructing a truth table.

Suppose Jonas is a philosophy student who is just beginning to learn formal logic. I present to him the sentence $(p \rightarrow (q \rightarrow p))$, which I'll call φ, and I ask whether φ is a tautology or not. Jonas isn't terribly confident in logic yet, and doesn't have a good sense of how weird conditionals like φ will come out—after all, it's not a logical form we encounter much in the wild.

So, how does he go about investigating my question? How does he conduct an inquiry into it? He doesn't go out and gather empirical evidence as, say, Jane does when she wants to know how many windows there are in the Chrysler building. Rather, he conducts logical reasoning. In particular, he produces the following truth table, one row at a time:

p	q	$\varphi = (p \rightarrow (q \rightarrow p))$
T	T	T
T	F	T
F	T	T
F	F	T

And then he notes that, on each possible truth assignment, the sentence is true, and so concludes it's a tautology.

In Hacking's framework, we can represent Jonas' reasoning as follows. There are sixteen ways the final column might run:

T	T	T	T	T	T	T	T	F	F	F	F	F	F	F	F
T	T	T	T	F	F	F	F	T	T	T	T	F	F	F	F
T	T	F	F	T	T	F	F	T	T	F	F	T	T	F	F
T	F	T	F	T	F	T	F	T	F	T	F	T	F	T	F

Of course, all but the first of these is logically impossible, but in Hacking's framework there are nonetheless possible worlds at which each is true. So let's say Jonas starts with equal credence in each of these possibilities. So he has credence $1/16$ that φ is a tautology. Now he completes the first row and that of course rules out all worlds in which the first row is false—that is, the right-hand half of the array above. He updates by conditionalizing on that, and comes to assigns credence $1/8$ to each of the eight remaining possibilities for the final column. Moreover, he comes to assign credence $1/8$ to the claim that φ is a tautology, which he knows is equivalent to the first column from the array above. Then he completes the second row, and that leads him to assign credence $1/4$ to the claim that φ is a tautology; completing the third row leads him to assign credence $1/2$ to that claim; and completing the fourth row leads him to be fully certain of it.

So, on this approach, logical reasoning is formally very close to empirical evidence-gathering. In both cases, we do something—reason or acquire evidence—that allows us to rule out certain possibilities. And then we update as Personal Sentential Conditionalization demands.

We can now use Hacking's approach to represent the situation of Jane the glazier as she is counting the windows of the Chrysler building and overhears her neighbour's conversation. She must choose between two ways of proceeding: on the first, she keeps counting the windows to gather new empirical evidence about one issue; on the second, she reasons from the evidence she already has about another issue. Both can be represented in the value of information framework once it is transposed into Hacking's framework. We might say that continuing to count the windows is the evidence-gathering episode that will teach Jane the true sentence of the form 'The Chrysler building has *n* windows', while reasoning from her evidence is the one that will teach her the true sentence of the form 'Park Avenue runs in direction *d*'. After she has learned the evidence from her neighbour that Madison Avenue runs NE-SW, we might suppose that she has equal credence in each of the following sentences:

'Park runs NE-SW & Madison runs NE-SW & Madison and Park are parallel'

'Park runs N-S & Madison runs NE-SW & Madison and Park are parallel'

'Park runs NW-SE & Madison runs NE-SW & Madison and Park are parallel'

'Park runs E-W & Madison runs NE-SW & Madison and Park are parallel'

That is, credence 1/4 in each. And we might suppose she knows that reasoning from her evidence will teach her exactly which of these is true.

And we might suppose that, since she has reached 2,000 in her window count by that point, and she knows antecedently that there are at most 10,000 windows, she has equal credence in each of the following sentences:

'The Chrysler building contains 2,001 windows'

'The Chrysler building contains 2,002 windows'

...

'The Chrysler building contains 9,998 windows'

'The Chrysler building contains 9,999 windows'

'The Chrysler building contains 10,000 windows'

That is, credence 1/8,000 in each. And we might suppose she knows that gathering her evidence will teach her exactly which is true.

So, we can see that Jane has all the ingredients she needs to calculate the expected pragmatic or epistemic or all-things-considered value of each of the two ways of proceeding: gathering further empirical evidence or reasoning from existing evidence. And so she can compare these and pick whichever has greater expected value. That is, just as we saw in the previous chapter, the general norm that guides Jane, and which replaces EP_a, EP_o, and ZIP, is simply the norm that exhorts us to maximize expected utility.

5.3 How should we direct our attention?

It's natural to describe Jane the glazier in front of the Chrysler building as choosing how to direct her attention: to the windows of the building or to the evidence from her surroundings. Like gathering evidence or reasoning from evidence, directing attention is simply an action we perform—we might call it an epistemic action, since it is undertaken at least partly in order to alter our epistemic state. Indeed, gathering evidence or extracting information from evidence are perhaps just two species of directing attention. And so, since it's an action we perform under uncertainty about its outcome, it is something that a totalizing version of Bayesianism will want to govern. And indeed we might naturally say that we should direct our attention in whichever way among those that are possible maximizes expected pragmatic or epistemic or all-things-considered value.

You might worry that we don't always choose how to direct our attention, and so it is not appropriate to represent us as making a decision between directing it in one way rather than another. But while its name suggests that decision theory is only about decisions we voluntarily and consciously make, and rational choice theory only about options we voluntarily and consciously choose, there is no need to treat them like that. We can just as well use their tools to assess choices that are unconscious, involuntary, or both. We can say of these choices too that they maximize or fail to maximize expected utility. And, while we might hesitate to praise someone for doing something unconsciously and involuntarily that does maximize expected utility or blame someone for doing something unconsciously and involuntarily that fails to so maximize, nonetheless, we can say that it's better to do the former and worse to do the latter, and that whatever unconscious process is driving the involuntary action is serving the person who has it well in the first case and poorly in the second. That is to say: decision theory is a way of evaluating, from a particular point of view that includes credences and utilities, the selection of an option from a menu. Nothing in the theory requires that this selection is performed by a person; it might be performed by a non-human animal, an organism, a computer, a state, or an algorithm. And nothing in it requires that it be done intentionally or consciously.

This point also helps the totalizing view to accommodate Georgi Gardiner's (2022) observation that often we assess an individual's patterns of attention over a period, rather than individual acts of directing attention. We notice that people tend to pay attention to this rather than that. And we might notice that, while each individual instance within this pattern might be considered permissible when considered in isolation, the pattern itself is not. Gardiner argues that we should therefore take a more virtue-theoretic approach to our evaluation of a person's patterns of attention. But I think we should rather use the Value of Information approach I've been describing, but where the option to be assessed is not just an individual act of directing attention in one way rather than another, but instead a disposition to direct attention that is manifested over a period of time. This allows us to assess Friedman's choice to continue directing her attention to the Chrysler windows, which is a single action; and it allows us to assess an individual's lifelong disposition to direct attention in particular ways, which is the sort of thing a virtue-theoretic approach will consider; but it also allows us to assess dispositions manifesting over shorter timescales than the virtue theorist will consider, such as when someone, during a week-long period of uncharacteristically heightened anxiety, focuses only on the possibilities that lead to mortal danger.

The Value of Information approach I've been proposing in this second part of the essay can also help us to understand the interesting phenomena of group attention that Gardiner discusses. She points out that groups can contain individuals who direct their attention in ways that are bad for them but end up helping the group, and there are groups in which all individuals direct their attention in good ways, but the group itself performs poorly. This recalls what Mayo-Wilson et al. (2011) call *the independence thesis* in the philosophy of science. That thesis says that epistemically rational individuals might form epistemically irrational groups and epistemically irrational individuals might form epistemically rational groups.

Suppose, for instance, that each individual in a group begins with the same prior credences. And suppose that each directs her attention to just one feature of a situation, leading her to collect evidence in a very biased fashion that is, in expectation and from her individual point of view, worse than not collecting it. But suppose further that each individual directs attention to a *different* single feature of the situation. Then the group itself might end up with a very balanced and rich set of evidence, since there will be no overlap between the pieces of evidence gathered by the individuals. In this situation, each individual will end up with credences that have lower epistemic and pragmatic utility in expectation than they would have had if she had directed her attention in a less focussed way and gathered evidence more evenly. But, if we take the group's credences to be the shared prior credences updated on the group's total evidence—which is this balanced and rich set—then they will have higher epistemic and pragmatic utility in expectation than they would have had if each individual had gathered the evidence more evenly and ended up with a total body of evidence that was also very balanced, but this time much less rich, since there would be a lot of overlap.

In general, the teleological approach embodied in the Value of Information framework I've been describing is useful in non-ideal epistemology—its totalizing ambitions are not restricted to ideal epistemology. In that approach to epistemology, we look at different components of our epistemic practices—how we inquire, how we reason, how we direct attention, and so on. Each of these practices leads ultimately to the credences we have. So we can evaluate them by the expected epistemic or pragmatic utility of the credences they'll lead to. Non-ideal epistemology recognises that we are limited, and so the range of epistemic practices we can engage in is limited—Friedman's window counter can attend to the windows or to what her evidence tells us about the direction Park Avenue, but she can't do both; as I sit on a short train journey, I have the time to draw out some inferences from the evidence I received last night chatting to a friend and I have the time to read about my latest interest on Wikipedia, but I can't do both. And a teleological approach is well suited to that as well: when you assess a particular practice we engage in, you compare it only with others that are actually available to limited creatures like us.

So you don't criticise the window counter for not attending to both sources of information, since this is simply beyond her capacities. While some of the main versions of this teleological approach to epistemology has often focused on general norms for ideal epistemic agents—betting arguments and accuracy arguments for Probabilism, Conditionalization, and so on—it is ideally suited to answering the sorts of questions that interest non-ideal epistemologists as well.

5.4 Staffel on personal probabilities

In her forthcoming book, Julia Staffel (ta) raises four problems for Hacking's personal probabilities framework, as well as my extension of it. In this section, I'll work through these and say why I don't find them as worrying as Staffel does.

The problem of logical obtuseness

The first problem Staffel raises only to present and accept what I take to be the correct response to it, and indeed the response that Hacking himself endorsed. The worry runs like this: Classical Bayesianism is clearly too demanding, since it is surely not irrational to be less than certain of some logical fact you've never had the time to establish for yourself. But Hacking's personal Bayesianism is surely not demanding enough. After all, if I haven't done the reasoning that would lead me to be certain of some very simple tautology, such as 'It's raining or it's not', then I am not irrational if I am less than certain of it. But this is to permit a certain obtuseness about logical truths. Surely there is simply something irrational about someone who is less than certain of this sentence?

It is in fact the central insight of Hacking's paper that the Value of Information framework, transposed to his Personal Sentential Bayesian framework, can help us to answer this objection. The point is this: many simple logical truths are extremely easy to establish and extremely useful to know, and so if you aren't certain of them, that means you haven't undertaken the logical reasoning required to become certain of them, but that logical reasoning is so low cost and so high gain that not undertaking it is irrational; it is, in expectation, better to undertake that reasoning than not to, and so someone who hasn't is irrational.

As an interesting side-note, I think Hacking's response here also explains why Jonas, our fledgling logician from above, might initially be uncertain of $(p \rightarrow (q \rightarrow p))$ without being irrational, even though this is a simple logical truth, and it is not difficult to establish. We don't judge him irrational for not being certain of it, because it is not a logical truth that is particularly useful for him to know. He will rarely find himself wondering whether he can appeal to it in his reasoning. And so, before the question was put to him by his logic teacher, it was not worth the effort required of him to establish it for himself, even though that effort would have been minimal.

The problem of acting on credences formed mid-reasoning

Sometimes, we embark on a series of evidence-gathering episodes or a series of pieces of logical reasoning that we hope we'll complete before we have to make a particular decision, but we get interrupted midway through and have to make the decision then and there. Which credences should we use when this happens? Our midway credences, that is, the ones we have when we are interrupted and must face the decision? Or the ones we had at the beginning of our investigation? Or some other credences?

Providing each step in our series of evidence-gathering episodes or pieces of logical reasoning maximizes expected epistemic or pragmatic or all-things-considered utility, we should use our midway credences, for while our credences at the beginning of the series would prefer most to have completed the series and used the credences we'd have had then, they nonetheless prefer using the midway credences to using themselves.

Staffel contends that, while this is the correct answer for a series of evidence-gathering episodes, it's the wrong answer for a series of pieces of logical reasoning. Intuitively, the person who acts on the basis of credences they've formed after an incomplete sequence of logical reasoning from their evidence is doing worse than the person who acts on the basis of credences they've formed after an incomplete empirical inquiry. Were Jane interrupted in her counting of the windows and asked to submit her bid for the glazing contract midway through, she'd have to choose how to cost it using her credences at that time. She'd be upset, of course, because she'd hoped to improve those credences more before having to do this. But we wouldn't judge her negatively: she'd be doing the rational thing in a suboptimal situation. On the other hand, were she interrupted after learning the direction of Madison Avenue, but before she'd had any chance to bring that evidence together with her evidence that Madison Avenue and Park Avenue are parallel, and offered a bet on the direction of Park Avenue, she'd again have to choose on the basis of credences formed before a full investigation is complete. And again she might be upset, and again we wouldn't necessarily judge her irrational, but in this case, Staffel thinks, we would judge her negatively in some sense: we'd think the situation is worse than in the previous version of the case.

I don't share this intuition myself. I do think we sometimes judge people more negatively in these sorts of cases, but that's only when we think that, upon learning that they must now make a decision, they would have time before the decision must be made in which to carry out the remaining logical reasoning. So, for instance, if there were enough time for Jane to reason to the direction of Park Avenue before having to decide whether to accept the bet offered, and she didn't, we'd of course judge her negatively. But the Value of Information approach predicts this. And if she didn't have the time—if the decision to accept or reject the bet must be made immediately—we shouldn't judge her negatively. To be rational is, roughly, to do the best with the resources you have. Of course, in some sense, Jane has the resources to acquire better credences about the direction of Park Avenue, for she has all the evidence she needs. But in the sense that matters, she doesn't have those resources, because she doesn't have the time to draw out the relevant consequence of her evidence.

The problem of learning logical falsehoods

In the Classical Bayesian setting, if you become certain of an empirical falsehood, you must update upon it by standard Bayesian conditionalization. This is established by betting arguments and accuracy arguments (Rescorla, 2020; Pettigrew, 2023b; Staffel & De Bona, 2023). These arguments transpose perfectly into Hacking's personal probability setting and entail that, if you become certain of a empirical falsehood *or a logical falsehood*, you're required to update upon it in the usual way. Staffel submits that this gives the right answer for empirical falsehoods, but the wrong answer for logical falsehoods.

Staffel is particularly concerned about cases in which we become certain of logical falsehoods because we make very obvious logical errors in our reasoning. Suppose Jonas makes a simple slip when producing the truth table for $(p \to (q \to p))$—he writes 'F' in the second row, for instance, and he thereby becomes certain of the sentence 'If p is true and q is false, then φ is false', which is false. According to Personal Sentential Bayesianism, he should update upon his new certainty in this falsehood by conditionalizing on it, and if he does so, he's fulfilled the requirements of rationality. And Staffel thinks this is wrong.

I agree that something has gone wrong here, and it's not only that Jonas has become certain of a falsehood. He has become certain of a falsehood when he really should not have. He has done something epistemically suboptimal—though it is unclear to me whether we'd call it irrational. But this suboptimality does not lie in his *response* to having become certain of this falsehood: he is rationally required to update by conditionalizing, given he becomes certain. The suboptimality lies in becoming certain of the falsehood in the first place. And we can trace that to the poor functioning of the certainty-forming method he used when completing the second row. It is akin to a case in which my visual system malfunctions briefly and I become certain the hat of the person in front of me is yellow when in fact it's green. In this case, I must respond to this new certainty by conditionalizing on it. And that will of course let my error reverberate through my credences. And this is suboptimal. And it is the same suboptimality that Jonas exhibits.

The problem of interpreting the priors

According to the standard Bayesian picture, we begin our epistemic life with what we might call our ur-prior credences. These are the ones we have before we gather any evidence. Then the evidence comes in and we update on it by conditionalizing.

It has become popular recently to interpret these ur-prior credences as encoding your epistemic or evidential standards (Schoenfield, 2014; Callahan &

Titelbaum, ta). If your ur-prior credence in a hypothesis conditional on some other proposition is greater than your unconditional ur-prior credence in that hypothesis, then you take the proposition to be evidence for that hypothesis—that's part of your evidential standards. Many people who take this view hold that there are many different permissible evidential standards: you and I might both agree that seeing a coin land heads eight times out of ten is some evidence that the coin is biased towards heads, but we might disagree on the strength of that evidence, with you taking it to be strong evidence and me taking it to be weak evidence.

Staffel's fourth concern about the personal Bayesian framework is that it does not allow us to interpret the ur-prior credences in this way. After all, our ur-prior credences will be uncertain of many logical truths. For instance, for Jane the glazier, learning that Madison Avenue and Park Avenue are parallel and that Madison Avenue runs NE-SW does not increase her credence that Park Avenue runs NE-SW. Her credence in the latter only increases after she does the logical reasoning required to realise that her evidence entails that Park Avenue runs in that direction. But surely we don't want to say that Jane takes the fact about the direction of Madison Avenue and the fact that Madison Avenue and Park Avenue are parallel to have no evidential bearing on hypotheses about the direction of Park Avenue? So personal Bayesianism doesn't seem to get this right.

My own view is that it's a mistake to think of ur-priors as encoding epistemic or evidential standards. They are simply our credences before evidence comes in. They are nothing more. They don't encode any judgments about evidential bearing, unless of course we assign credences to sentences that explicitly talk about evidential bearing.

Indeed, there's good reason to think that they don't encode such attitudes. Let's suppose, for instance, that we have someone unsure about the evidential bearing of some piece of evidence on some hypothesis. They are certain that it does have a bearing, but they're uncertain what bearing it has. They think it equally likely that it will tell against the hypothesis to a certain extent as they are that it will tell in favour of the hypothesis to that same extent. Then, upon learning the evidence, this person will keep their credence in the hypothesis exactly the same—they will neither increase it nor decrease it. But that doesn't mean they judge that the evidence has no bearing—indeed, we've specified they know it does! The point is that you cannot read off judgments about evidential dependencies from dependencies between credences or the lack of such dependencies.

Chapter 6

The Puzzle of Unwanted Evidence

So far, we have been concerned only to ask when we should gather evidence, and which evidence we should gather when we do. We have not considered what we should do with the credences we form after updating on the evidence we acquire. You might think that this is easily answered: we should use these credences as a guide to action and as a starting point for future reasoning. These credences form your point of view now, and it is the point of view from which you should choose what to do and what to think.

However, there are cases when that isn't obviously the right thing to do. In this chapter, I'll ask when that's the case and what we should do instead in those cases. In the end, I don't have a satisfactory answer to offer yet, so I will content myself with describing the problem and explaining why certain natural responses don't work.

6.1 The puzzle

First, some cases in which the problem arises.

- You might decide to gather evidence because you're pretty confident you'll face one decision problem and, relative to that, gathering this evidence is better in expectation than not gathering it, from the point of view of your priors; but, after gathering the evidence, you realise you'll actually face a different decision problem and, relative to that, gathering the evidence was not better in expectation than not gathering it, from the point of view of your priors. Facing the decision you'll actually face, which credences should you use? The priors or the ones you have on the basis of the evidence you gathered?

  Example. Recall the case of the unmarked clock from above (Example C, Section 2.6). Suppose you start with equal credences over 1 and 2 and ... and 12. If it's 2, you'll learn it's either 1 or 2 or 3; if it's 3, you'll learn it's either 2 or 3 or 4; and so on. At first you think you are going to face the choice between *Low*, which will pay a million dollars if the clock is pointing to 1 or 2 or . . . or 6, and *High*, which will pay a million dollars if the clock is pointing to 7 or 8 or . . . or 12. Your priors expect gathering this evidence to be better than not gathering it. But then, after gathering it, you learn you'll in fact face the choice between *Odd*, which will pay a million dollars if the clock is pointing to an odd number, *Even*, which will pay a million dollars if the clock is pointing to an even number, and *Neither*, which will pay nothing for sure. Now those priors expect gathering the evidence to be worse in expectation than not gathering. But it's too late. You've gathered it. So how should you choose?

- You might be quite right about *which* decision you'll face, but wrong about when. Thinking the decision will be further in the future, you might embark on a sequence of two evidence-gathering episodes with the following feature: the first on its own is worse in expectation than not gathering that evidence, relative to your priors, but the first and second taken together are better in expectation than not gathering them both. Thinking you have time to complete both, you embark on the first. But, you are then made to face the decision without being able to carry out the second evidence-gathering episode. Facing this decision earlier than you expected, which credences should you use? The priors or the ones you have on the basis of the evidence you have been able to gather so far?

  Example. You have a headscarf in your pocket that is one of four colours: red, rose, peach, or orange. These colours are not very easily distinguishable. I have the opportunity to observe the headscarf under two different lighting conditions. Under the first, if it's red or rose, I'll learn it's red or rose or peach; and if it's peach or orange, I'll learn it's rose or peach or orange. Under the second, if it's red or rose, I'll learn it's red or rose or orange; if it's peach or orange, I'll learn it's red or peach or orange. But, importantly, in both cases, I won't also learn that I learn these propositions; my evidence will not be luminous to me.

  Now, suppose I begin with uniform credences over the four colours—1/4 each. And suppose I know I'll face a bet on whether the headscarf is rose or peach: it will pay me £4 if it is and it will lose me £6 if it is not. My priors will reject this bet, since it will lose me £1 in expectation. After viewing the headscarf under the first lighting condition, however, my posteriors will accept the bet—whatever I learn, my credence that the headscarf is rose or peach will be 2/3, and so the bet will earn me £2/3 in expectation. So my priors expect gathering the evidence to be worse than not gathering, in expectation. Nonetheless, I gather because I think I'll also be able to view the necktie under the second lighting condition, and viewing under both conditions is better in expectation than viewing under neither. But I get interrupted and must choose after seeing it only under the first lighting condition. So how should I choose?

- I have talked so far as if all our evidence comes from evidence-gathering episodes we choose to undertake. But of course most of our evidence comes to us without us choosing to gather it. Jane the glazier, for instance, learned the direction of Madison Avenue without choosing to do so; I learn the colour of the hat the person in front of me is wearing without intentionally gathering that evidence; and so on. So now suppose that you inadvertently gather some evidence and your prior thinks that gathering that evidence is worse in expectation than not gathering it, given a decision you will face. Facing this decision having gathered the evidence in question, which credences should you use? The priors or the ones you have on the basis of the evidence you inadvertently received?

Example. A version of the coloured headscarf example works here. Suppose viewing it under the second lighting condition is not a possibility—I can only view it under the first lighting condition. So I'd prefer not to. But nonetheless, inadvertently, I do see it under that condition. And now I face the bet on peach or rose. How should I choose?

In each of these cases, you receive what you might call *unwanted evidence*. In two of the cases, you did want it at the beginning, but that was because you had a faulty understanding of the situtation: you thought you'd face a different decision problem with your posterior credences, or you thought you'd face the decision at a different time. Once you learn which decision you'll face and when, your priors at least would not want the evidence you've in fact received.

6.2 The Principle of Total Evidence solution

A natural response to the problem of unwanted evidence is to say that you should choose using your posterior credences. After all, they are based on a greater body of evidence. If you choose using your priors instead, you violate the so-called Principle of Total Evidence, which says that your credences should be based on all of your evidence, not only some part of it. The problem with this response is two-fold. First, Hosiasson's central purpose in devising the Value of Information framework and proving the central theorem was precisely to establish the Principle of Total Evidence. The title of her paper is this: 'Why do we prefer probabilities relative to many data?'. Why do we prefer credences based on more evidence rather than less? So, in cases in which you would choose not to gather the evidence, it seems we can't appeal to the Principle of Total Evidence to argue that you should use your posteriors—that principle is only true because, relative to the sorts of evidence-gathering episodes we usually consider, you would choose to gather the evidence.

The second problem with this proposal is that there are clear cases in which we should not use our posteriors. The unmarked clock provides an example. Suppose I know I will face the choice between *Odd*, *Even*, and *Neither*. My priors would prefer I didn't look at the clock before choosing. But suppose that

I am nonetheless shown the clock and receive the evidence it gives me. It is clear I shouldn't use the credences I have. After all, using them is guaranteed to lose me money: if the clock points at an odd number, my posterior credence it points at an even number is 2/3, and so I'll pick *Even* and lose; if the clock points at an even number, my posterior credence it points at an odd number is 2/3, so I'll pick *Odd* and lose. So, in this case, my posterior is such a bad guide to action that this is visible from any point of view—not just the point of view of my priors. Indeed, if we were to ask my posteriors, would you use those posteriors, they'd say no.

6.3 The Self-Recommending Solution

This suggests a second response to the problem of unwanted evidence. Perhaps we should test priors and posteriors for whether they recommend using themselves. There are four possible situations:

(I) Both priors and posteriors agree that using the posteriors, whatever they may be, is best. In this case, you're required to use the posteriors. Standard cases of learning factive and partitional evidence are like this. By the Value of Information theorem, the priors will think the posteriors are best. And because the posteriors place all of their credence on possibilities at which the posteriors you'll have are the posteriors themselves, they'll expect using the posteriors to be best.

(II) Prior and posterior agree that using the prior is best. In this case, you're required to use the prior. As we just saw above, Unmarked Clock is such a case.

(III) Prior thinks prior is best, and posterior thinks posterior is best. In this case, you're permitted to use either prior or posterior. Viewing the coloured headscarf under the first lighting condition is such a case. The prior will reject the bet, while each of the possible posteriors will accept it. So the prior will prefer using itself to choose and the posterior will prefer using the posteriors, whatever they may be.

(IV) Prior thinks posterior is best, and posterior thinks prior is best. We haven't met a case like this, but they do exist.

The problem with this suggestion, which we might call *the self-recommending solution*, is that there are cases in which it gets clearly the wrong answer. Consider an instance of evidence-gathering that you know in advance might be deceptive. There are three possibilities: if you are in the first, you'll learn you are; if you're in second, you'll learn you are; if you are in the third, you'll learn you're either in the first or second. Now consider a bet that you're in the first or second possibility, which will pay you £1 if you are, but will lose you £9 if you aren't. And suppose you have uniform priors over the three possibilities. Then your priors will reject this bet, but each of your possible posteriors will accept it. So your prior thinks your prior is best. However, each of your posteriors thinks you are best using your posteriors, whatever they are. After all, each of them assigns zero credence to the third possibility, which is where you are deceived. But it seems that, if I were to be in the third possibility, with posteriors that are certain I'm in the first or second possibility, I shouldn't use them. So I think the self-recommending solution fails.

6.4 The Fragmented Evidence Solution

Daniel Greco (2019) proposes a solution to a related problem about higher-order evidence. He suggests that, in cases like the unmarked clock, not all of our evidence is available for all of our choices. Perhaps some of our evidence is available for choices made at the subpersonal level by the visuomotor system: perhaps, if the clock in fact points at 2, the visuomotor system has access to the evidence that it points and 1 or 2 or 3; so, if you are asked to point at where the hand was pointing, you'd use that evidence and perhaps always point somewhere between 1 and 3. But perhaps other evidence is available for choices made at the personal level when you come to assert something, for instance: perhaps, if the clock in fact points at 2, the system that leads to assertions has access to different evidence, such as your best guess as to where the hand points. So, if you are asked where the hand points, you'll assert something based on this other evidence.

The problem with this response is that it doesn't seem to help when we adapt our case of the unmarked clock. Suppose that, in order to indicate how you'd like to choose in the decision between *Odd*, *Even*, and *Neither*, you are asked to point to a piece of paper that lists the odd numbers, a piece of paper that lists the even numbers, and a piece of paper that lists no numbers. It is therefore your visuomotor system that carries out this choice. And it has access to the evidence that the hand points at 1 or 2 or 3. But that simply leads us back to our original question: should the visuomotor system appeal to the priors or to the posteriors based on that evidence? The fact that different systems have access to different bodies of evidence doesn't settle whether they should use the credences based on them or not.

Chapter 7

Why should we not resist evidence?

In this final chapter, I want to ask what the Value of Information approach might say about what Mona Simion (2023) calls 'epistemic duties to believe', and particularly what it says about the sorts of violations of those norms that she gathers together under the heading of 'resistance to evidence'. The core of Simion's concern is that, in the past, epistemologists have based their assessment of an individual's doxastic state—the justification or rationality of a belief or credence—entirely on the evidence the individual in fact has, rather than basing it on both the evidence they have *and the evidence they should have had*. This means that the racist who simply resists any evidence that undermines their racist beliefs will count as rational and justified, as will the sexist who ignores evidence provided by a woman, or the climate denier who simply does not take on evidence contrary to their position. Simion seeks an epistemic duty that requires us not only to believe when we have collected evidence that supports a proposition, but also to believe when there was evidence at hand that supported that proposition, whether or not we in fact collected it. Here's the norm she gives:

> **DTB:** A subject S has an epistemic duty to form a belief that p if there is sufficient and undefeated evidence for S supporting p.

Of course, much is going to turn on what it means to say that *there is* sufficient and undefeated evidence, and Simion gives a detailed account of this. Using the Value of Information theorem, we might offer an alternative account: *there is* evidence available to an individual when the cost of gathering it would be very small, and certainly greatly outweighed by the expected utility of gathering it. (One hiccup here is that there might be very very many different pieces of evidence available where the cost of gathering each piece is very small, but we can't gather it all. But let's bracket those cases; the ones we consider here are not like that.)

So now let's consider the sort of case Simion has in mind and see what the Value of Information theorem tells us about them.

> **Case 1: Testimonial Injustice.** Anna is an extremely reliable testifier and an expert in the geography of Glasgow. She tells George that Glasgow Central is to the right. George believes women are not to be trusted, and therefore fails to form the corresponding belief. (Simion, 2023)

I think the subjective Bayesian's assessment of this case is a little different from Simion's, since they deal with the agent's subjective prior credences, while Simion works with a notion of evidential probability that many subjective Bayesians—and certainly this particular subjective Bayesian!—disavow. On perhaps the most natural subjective Bayesian reading, the case of George and Anna isn't a case of resistance to evidence, but rather a case of irrational priors. After all, let's take the evidence that George obtains in this situation to be that Anna says Glasgow Central is to the right. He might well incorporate that evidence exactly as the Bayesian says he should and yet retain a low or middling credence that Glasgow Central is to the right. For Simion says that George believes women aren't to be trusted, and so this is something that is encoded in the credence function he has when he meets Anna and hears her testimony. The Bayesian says he should conditionalize on his priors, but doing so will lead him to have something pretty close to his previous middling credence about the direction of Glasgow Central, since he'll think Anna's testimony is not much better than chance as an indicator of the truth. For the Bayesian, the situation is structurally akin to a case in which I irrationally believe that the thermometer on my wall is completely broken when in fact it's very accurate, and so when it tells me that the temperature is 20C, I update on that evidence exactly as the subjective Bayesian says I should, conditionalizing my priors on it, but it doesn't shift my credences, because my priors were irrationally inaccurate and treated the thermometer's reading are almost independent of the true temperature.

So, for the subjective Bayesian, George is certainly flawed, but it's not because he is resistant to the evidence Anna gives him; or, at least, it isn't because he's resistant to the evidence in the sense that he fails to incorporate it. It is rather because he has an irrational prior that leads him to have an irrational posterior after he does incorporate the evidence in the way his prior demands.

Of course, his irrational prior might be the result of having resisted evidence in the past. There are at least two ways George might have ended up with that prior:

On the first, his ur-prior, the credence function he has at the beginning of his epistemic life, might have assigned very low credence to the reliability of women's testimony, and that will be judged irrational since it's taking an extreme stand on a proposition about which George had no evidence at that time. What's more, if he assigns higher credence to the reliability of men's testimony, say, we will judge it further irrational because it differentiates between two cases when he has no evidence to justify such differential treatment.

On the second way he might have arrived at his irrational prior, his ur-prior might have assigned middling credence to the reliability of women's testimony, just as it did to the reliability of everyone else's testimony, but then as he went through life he incorporated any evidence he received that told against women's reliability and failed to incorporate any evidence he received in its favour. This leaves him with the biased credence function he has when he meets Anna and hears her testimony. In this case, George exhibits genuine resistance to evidence he received, and Oddie's version of the Value of Information theorem tells us what went wrong with him: he failed to incorporate evidence when incorporating it would have improved his epistemic situation in expectation.

Let's turn now to a case raised by Simion in conversation:

> **Case 2: Climate change denier** Jon denies that there is an anthropogenic component to current dramatic changes in Earth's climate. Over the years, this has become such a large part of Jon's thinking that it constitutes part of his identity. A great deal of evidence to the contrary is available to him, but he resists it, perhaps unconsciously, because to face it and incorporate it properly would be to lose a belief that forms part of who he is; losing that belief would be very costly to Jon, resulting in anguish, disorientation, and alienating him from the epistemic bubble into which this belief has drawn him.

Surely, Simion contends, the Value of Information approach says that, in this case, Jon should not incorporate the evidence he has; and just as surely this is the wrong answer. Given the pain it will cause Jon to lose his belief that current climate change is entirely naturally caused, it almost certainly outweighs any expected pragmatic or epistemic utility he'll gain by gathering the evidence that will lead to this. So, even from his current point of view, where he assigns very high credence to the proposition that climate change is naturally caused, and therefore very low credence to the evidence he might gather changing his mind, the negative effects of changing his mind are so great that he still gives higher expected utility to not gathering the evidence. So, the Value of Information theorem says, he does nothing wrong by not gathering it. And that, Simion contends, is the wrong answer.

I think this is a case where it is helpful to distinguish two ways in which we might evaluate someone's actions: a purely subjective one, and a slightly more objective one. In the purely subjective sense, Jon indeed does nothing wrong: he best serves the ends and values that he actually has by avoiding the evidence that might change his mind about climate change; relative to his actual credences and actual utilities, this is what maximizes expected value. But in the more objective sense, we might criticize exactly those ends and values and so say that what he ought to do is not what maximizes in expectation the utilities that encode them. We might say that Jon ought not to have those values and the utilities that encode them; he ought not to have built his identity around an empirical belief in that way when doing so would make him so resistant to learning anything that would unseat it. If we do this, we might instead evaluate him not from the point of view of his actual credences and actual utilities but from the point of view of his actual credences and the utilities *he should have instead*—utilities that do not place such negative value on coming to believe that climate change is anthropogenically caused. And once we do that, the Value of Information theorem delivers the result we want: Jon should gather the evidence.

Recognising these different ways in which we might evaluate someone's evidence-gathering behaviour reveals a flexibility in the Value of Information approach I haven't had cause to highlight so far. Formally, the approach needs a vantage point from which to assess an evidence-gathering episode, a representation of that evidence-gathering episode as an evidence function, and an account of how the individual would respond to the evidence should they receive it. In the pragmatic case, the vantage point consists in a probability function, a decision problem they will face, and a utility function: the probability function is typically taken to be the actual credences of the would-be evidence-gatherer at the point at which they have to decide whether or not to gather the evidence, and their utility function encodes their actual values at that point. And in the epistemic case, the vantage point consists in a probability function and an epistemic utility function. As the example of Jon shows, in the pragmatic case, sometimes we might want to use other utilities than the individual's own; perhaps their own are immoral or unreasonable or self-undermining or in some other way flawed. But equally we might want to use other probabilities than the individual's own credences. We might want to use evidential probabilities, if we think there are such things. Or we might want to use the credences the individual *ought to have had*, not the ones they actually have, where we take these to be the ones they'd have had had they acted correctly in the past. This is another interesting theme that emerges in Simion's work on resistance to evidence.

Think again of Jon, the climate science denialist. We said he decides not to gather evidence that might persuade him that climate chance is manmade because that risks something he values greatly, even though he shouldn't. But consider his friend, Jim, also a climate denialist, who faces the same decision whether to gather evidence that might overturn his scepticism. Unlike Jon, Jim's utilities are in order: he ascribes no value to being a denialist and cares only about discovering the truth. However, because of how he's directed his attention in the past—or had his attention directed by his social environment—he has only ever picked up on evidence against the anthropogenic origins of climate change and he's always ignored or missed evidence in its favour, even when it's been readily available at a low cost to him. Let us stipulate that, in these past cases, the Value of Information approach would judge that Jim acted irrationally in not gathering this latter body of evidence. His current credences, shaped by his past history of gathering and not gathering different pieces of evidence, assigns such a high credence to the natural origin of climate change, and so thinks it so unlikely that gathering further evidence will change his mind that it's rational *from that point of view* not to pay the costs of gathering it. Nonetheless, Simion thinks and I agree there is a sense in which we want to say that Jim should gather it, just as Jon should. In this case, it isn't Jim's utilities we want to fix up, but his credences. It's not that we want to appeal to some objectively correct alternative credences, such as the evidential probabilities, but rather to the credences that Jim would have had had he gathered evidence in the past in the ways that were rationally required by his own credences at the time, but which he didn't gather.

Chapter 8

Conclusion

This brings us to the end of our exploration of the value of information and its relationship to the epistemology of individual inquiry. What, then, is the view? As noted at the end of the previous chapter, the Value of Information framework that grows out of Hosiasson's initial insight centres on a way of evaluating and comparing different evidence-gathering episodes with one another, with the choice not to gather evidence, and with other possible courses of action; and it provides a way of doing this evaluation from a pragmatic point of view, an epistemic point of view, and a combination of the two, which I've been calling the all-things-considered point of view.

The ingredients of this evaluative framework are as follows:

(I) *the vantage point from which the evaluation takes place*.

This includes:

(a) *a probability function*.

Typically, this is the would-be evidence-gatherer's credences when they must decide what to do; but equally it could be the evidential probabilities, if such exist, or the credences the would-be evidence-gatherer would have had had they behaved fully rationally in the past.

(b) *either*

(i) *a set of decision problems and a pragmatic utility function*.
The decision problems are those the individual thinks they might face with their credences, and the pragmatic utility function encodes their values and ends. These ingredients are needed for the pragmatic evaluation of evidence-gathering episodes.

or

(ii) *an epistemic utility function*.
This is needed for the epistemic evaluation of evidence-gathering episodes.

(II) *the available options*.

These might include evidence-gathering episodes represented by evidence functions, including the trivial evidence function, which represents not gathering evidence at all, but they might also include non-epistemic actions, such as making a sandwich.

(III) *the updating plan*.

This describes how the individual will respond to the evidence.

The crucial insight of the approach that descends from Hosiasson's original insight is that an evidence-gathering episode is an action like any other. She, Blackwell, and Good saw how to assess it for its pragmatic value, and Oddie saw how to evaluate it for its pure epistemic value. This allows us to bring such decisions within the ambit of the totalizing vision of rational choice theory, which seeks to govern at least the rationality of all of our actions.

The norms of evidence-gathering and its extended pursuit, which we call inquiry, are then simply instances of the more general norms of rational choice. If the correct theory of decision is expected utility theory, then these are:

Pragmatic norm of inquiry Gather evidence when doing so maximizes expected pragmatic utility.

Epistemic norm of inquiry Gather evidence when doing so maximizes expected epistemic utility.[1]

All-things-considered norm of inquiry Gather evidence when doing so maximizes expected all-things-considered utility.

But the framework is flexible enough that it will serve if a different theory of decision is correct, such as one that permits sensitivity to risk or ambiguity or both (recall Section 2.5).

What's more, by expanding our understanding of the possibilities over which credences and utilities are defined in our decision-making model so that they include impossibilities as well as possibilities, and by representing logical or *a priori* reasoning as ruling out such impossibilities, just as empirical evidence rules out empirical possibilities, we can understand logical and *a priori* inquiry using the Value of Information framework, and thereby provide a unified account of all inquiry, empirical and otherwise.

[1] Flores & Woodard (2023) ask whether there are genuinely epistemic norms on evidence-gathering. This seems a candidate for this. In the end, I don't think too much hangs on how we categorize norms into the pragmatic and the epistemic. The norm just given is epistemic in the sense that it evaluates something for how it serves purely epistemic ends. But the thing evaluated is not itself purely epistemic: it is not a credence or a belief, but rather an action. The action is of course aimed at gathering evidence, which in turn is aimed at changing a purely epistemic state, such as a credence or belief. Is that sufficient to make it epistemic? As I say, I'm not sure we need adjudicate this. The insights of Flores and Woodard's paper, which describes epistemic ways in which we criticize evidence-gathering episodes that violate certain norms, stand whether or not they support the claim that the norms themselves are epistemic.

Bibliography

Allais, M. (1953). Le comportement de l'homme rationnel devant le risque: critique des postulats et axiomes de l'école Américaine. *Econometrica*, *21*(4), 503–546.

Archer, A. (2021). The Aim of Inquiry. *Disputatio*, *13*(61), 95–119.

Berger, J. O. (1985). *Statistical Decision Theory and Bayesian Analysis*. New York: Springer.

Blackwell, D. (1951). Comparison of Experiments. *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, *2*, 93–102.

Bottomley, C., & Williamson, T. L. (ta). Rational risk-aversion: Good things come to those who weight. *Philosophy and Phenomenological Research*.

Bright, L. K. (2017). On Fraud. *Philosophical Studies*, *174*(2), 291–310.

Brown, P. M. (1976). Conditionalization and expected utility. *Philosophy of Science*, *43*(3), 415–419.

Buchak, L. (2013). *Risk and Rationality*. Oxford, UK: Oxford University Press.

Callahan, L. F., & Titelbaum, M. G. (ta). What are Epistemic Standards? In S. Goldberg, & M. Walker (Eds.) *Attitude in Philosophy*. Oxford: Oxford University Press.

Carr, J. (2021). A Modesty Proposal. *Synthese*, *198*(4), 3581–3601.

Das, N. (2023). The Value of Biased Information. *The British Journal for the Philosophy of Science*, *74*(1), 25–55.

de Finetti, B. (1931). Sul significato soggettivo della probabilita. *Fundamenta Mathematicae*, *17*, 298–329.

Dorst, K. (2020). Evidence: A Guide for Uncertain. *Philosophy and Phenomenological Research*, *100*(3), 489–757.

Dorst, K., Levinstein, B. A., Salow, B., Husic, B. E., & Fitelson, B. (2021). Deference Done Better. *Philosophical Perspectives*, *35*, 99–150.

Ellsberg, D. (1961). Risk, Ambiguity, and the Savage Axioms. *Quarterly Journal of Economics*, *75*(4), 643–69.

Falbo, A. (2023). Should epistemology take the zetetic turn? *Philosophical Studies*, *180*(10-11), 2977–3002.

Flores, C., & Woodard, E. (2023). Epistemic Norms on Evidence-Gathering. *Philosophical Studies*, *180*, 2547–2571.

Friedman, J. (2020). The Epistemic and the Zetetic. *Philosophical Review*, *129*(4), 501–536.

Gallow, J. D. (2021). Updating for Externalists. *Noûs*, *55*(3), 487–516.

Gardiner, G. (2022). Attunement: On the Cognitive Virtues of Attention. In M. Alfano, C. Klein, & J. de Ridder (Eds.) *Social Virtue Epistemology*. New York: Routledge.

Geanakoplos, J. (1989). Game Theory without Partitions, and Applications to Speculation and Consensus. Tech. rep., Cowles Foundation Discussion Paper.

Gibbard, A., & Harper, W. (1978). Counterfactuals and Two Kinds of Expected Utility. In E. F. M. Clifford Alan Hooker, James L. Leach (Ed.) *Foundations and Applications of Decision Theory*, vol. 13a of *University of Western Ontario Series in Philosophy of Science*, (pp. 125–162). D. Reidel.

Good, I. J. (1967). On the Principle of Total Evidence. *The British Journal for the Philosophy of Science*, *17*, 319–322.

Greaves, H., & Wallace, D. (2006). Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility. *Mind*, *115*(459), 607–632.

Greco, D. (2019). Fragmentation and Higher-Order Evidence. In M. Skipper, & A. Steglich-Petersen (Eds.) *Higher-Order Evidence: New Essays*, (pp. 84–104). Oxford: Oxford University Press.

Hacking, I. (1967). Slightly more realistic personal probability. *Philosophy of Science*, *34*(4), 311–325.

Hookway, C. (2006). Epistemology and Inquiry: The Primacy of Practice. In S. Hetherington (Ed.) *Epistemology Futures*, (pp. 95–110). Oxford University Press.

Hosiasson, J. (1931). Why do we prefer probabilities relative to many data? *Mind*, *40*(157), 23–36.

Huttegger, S. M. (2013). In Defense of Reflection. *Philosophy of Science*, *80*(3), 413–433.

Huttegger, S. M. (2014). Learning experiences and the value of knowledge. *Philosophical Studies*, *171*(2), 279–288.

Huttegger, S. M., & Nielsen, M. (2020). Generalizing Learning and Conditional Expectation. *Philosophy of Science*, *87*(5), 868–883.

Isaacs, Y., & Russell, J. S. (2023). Updating without Evidence. *Noûs*, *57*(3), 576–599.

Jeffrey, R. C. (1983). *The Logic of Decision*. Chicago and London: University of Chicago Press, 2nd ed.

Joyce, J. M. (1998). A Nonpragmatic Vindication of Probabilism. *Philosophy of Science*, *65*(4), 575–603.

Joyce, J. M. (2009). Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief. In F. Huber, & C. Schmidt-Petri (Eds.) *Degrees of Belief*. Dordrecht and Heidelberg: Springer.

Kelp, C. (2021). *Inquiry, Knowledge, and Understanding*. Oxford, UK: Oxford University Press.

Levinstein, B. A. (2023). Accuracy, Deference, and Chance. *Philosophical Review*, *132*(1), 43–87.

Lewis, D. (1981). Causal Decision Theory. *Australasian Journal of Philosophy*, *59*, 5–30.

Mayo-Wilson, C., Zollman, K. J. S., & Danks, D. (2011). The Independence Thesis: When Individual and Social Epistemology Diverge. *Philosophy of Science*, *78*(4), 653–77.

Myrvold, W. C. (2012). Epistemic values and the value of learning. *Synthese*, *187*(2), 547–568.

Nielsen, M. (2024). Only CDT values knowledge. *Analysis*, *84*(1), 67–82.

O'Connor, C., & Weatherall, J. O. (2018). *The Misinformation Age: How False Beliefs Spread*. New Haven: Yale University Press.

Oddie, G. (1997). Conditionalization, Cogency, and Cognitive Value. *British Journal for the Philosophy of Science*, *48*, 533–41.

Pérez Carballo, A. (2018). Good Questions. In J. Dunn, & K. Ahlstrom-Vij (Eds.) *Epistemic Consequentialism*. Oxford, UK: Oxford University Press.

Pettigrew, R. (2016). *Accuracy and the Laws of Credence*. Oxford: Oxford University Press.

Pettigrew, R. (2020a). *Dutch Book Arguments*. Elements in Decision Theory and Philosophy. Cambridge, UK: Cambridge University Press.

Pettigrew, R. (2020b). Logical Ignorance and Logical Learning. *Synthese*.

Pettigrew, R. (2023a). Bayesian updating when what you learn might be false. *Erkenntnis*, *88*, 309–324.

Pettigrew, R. (2023b). Bayesian Updating When What You Learn Might Be False. *Erkenntnis*, *88*, 309–324.

Pettigrew, R. (2023c). Epistemic Utility Arguments for Epistemic Norms. In E. N. Zalta (Ed.) *Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University.

Predd, J., Seiringer, R., Lieb, E. H., Osherson, D., Poor, V., & Kulkarni, S. (2009). Probabilistic Coherence and Proper Scoring Rules. *IEEE Transactions of Information Theory*, *55*(10), 4786–4792.

Quiggin, J. (1982). A Theory of Anticipated Utility. *Journal of Economic Behavior and Organization*, *3*, 323–43.

Quiggin, J. (1993). *Generalized Expected Utility Theory: The Rank-Dependent Model*. Kluwer Academic Publishers.

Ramsey, F. P. (1926 [1931]). Truth and Probability. In R. B. Braithwaite (Ed.) *The Foundations of Mathematics and Other Logical Essays*, chap. VII, (pp. 156–198). London: Kegan, Paul, Trench, Trubner & Co.

Ramsey, F. P. (1990). Weight or the Value of Knowledge. *The British Journal for the Philosophy of Science*, *41*, 1–4.

Rescorla, M. (2020). An Improved Dutch Book Theorem for Conditionalization. *Erkenntnis*.

Rosa, L. (2025). *Rules for the Inquiring Mind: A Unified Framework of Norms of Inquiry*. London, UK: Routledge.

Rosenstock, S., Bruner, J., & O'Connor, C. (2017). In Epistemic Networks, Is Less Really More? *Philosophy of Science*, *64*(2), 234–252.

Salow, B. (2018). The externalist's guide to fishing for compliments. *Mind*, *127*, 691–728.

Savage, L. J. (1954). *The Foundations of Statistics*. John Wiley & Sons.

Schliesser, E. (ta). Synthetic Philosophy: a Restatement. *Proceedings of the Aristotelian Society*.

Schoenfield, M. (2014). Permission to Believe: Why Permissivism Is True and What It Tells Us About Irrelevant Influences on Belief. *Noûs*, *48*(2), 193–218.

Schoenfield, M. (2017). Conditionalization does not (in general) Maximize Expected Accuracy. *Mind*, *126*(504), 1155–87.

Schultheis, G. (ta). Accurate Updating. *Philosophy of Science*.

Simion, M. (2023). Resistance to Evidence and the Duty to Believe. *Philosophy and Phenomenological Research*.

Skyrms, B. (1982). Causal Decision Theory. *The Journal of Philosophy*, *79*(11), 695–711.

Skyrms, B. (1990a). *The Dynamics of Rational Deliberation*. Cambridge, Mass.: Harvard University Press.

Skyrms, B. (1990b). The Value of Knowledge. *Minnesota Studies in the Philosophy of Science*, *14*, 245–66.

Staffel, J. (ta). *Unfinished Business: Rational Attitudes in Reasoning*. Oxford, UK: Oxford University Press.

Staffel, J., & De Bona, G. (2023). An Improved Argument for Superconditionalization. *Erkenntnis*.

Thorstad, D. (2022). There are no epistemic norms of inquiry. *Synthese*, *200*(410).

Torsell, C. (2024). Janina Hosiasson and the value of evidence. *Studies in the History and Philosophy of Science*, *106*, 31–36.

van Fraassen, B. C. (1984). Belief and the Will. *Journal of Philosophy*, *81*, 235–56.

van Fraassen, B. C. (1995). Belief and the Problem of Ulysses and the Sirens. *Philosophical Studies*, *77*(1), 7–37.

van Fraassen, B. C. (1999). Conditionalization, A New Argument For. *Topoi*, *18*(2), 93–96.

Vineberg, S. (2016). Dutch Book Arguments. In E. N. Zalta (Ed.) *Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University.

Willard-Kyle, C. (forthcoming). The Knowledge Norm of Inquiry. *The Journal of Philosophy*.

Williams, J. R. G. (2018). Rational Illogicality. *Australasian Journal of Philosophy*, *96*(1), 127–141.

Williams, J. R. G., & Pettigrew, R. (2023). Consequences of Calibration. *The British Journal for the Philosophy of Science*.

Williamson, T. (2013). Gettier Cases in Epistemic Logic. *Inquiry: An Interdisciplinary Journal of Philosophy*, *56*(1), 1–14.

Williamson, T. (2014). Very Improbably Knowing. *Erkenntnis*, *79*(5), 971–999.

Woodard, E. (2022). Why Double-Check? *Episteme*, (pp. 1–24).

Zollman, K. J. S. (2007). The communication structure of epistemic communities. *Philosophy of Science*, *74*(5), 574–587.

Zollman, K. J. S. (2010). The Epistemic Benefit of Transient Diversity. *Erkenntnis*, *72*, 17–35.