

Taking a Good look at the norms of gathering and responding to evidence*

Richard Pettigrew
Richard.Pettigrew@bristol.ac.uk

September 1, 2023

Abstract

In the recent philosophical literature on inquiry, epistemologists point out that their subject has often begun at the point at which you already have your evidence and then focussed on identifying the beliefs for which that evidence provides justification. But we are not mere passive recipients of evidence. While some comes to us unbidden, we often actively collect it. This has long been recognised, but typically epistemologists have taken the norms that govern inquiry to be practical, not epistemic. The recent literature challenges this assumption and uncovers a rich range of questions about the epistemic normativity of inquiry. In this paper, I approach these questions from the formal side of epistemology. Developing out of the philosophy of science, as it did, this branch of epistemology has long discussed inquiry. And, building on the insights of David Blackwell (1951) and I. J. Good (1967), it has produced a reasonably well-developed framework in which to understand norms of inquiry, both epistemic and practical. In the first half of the paper, I will present the pragmatic versions of this framework due to Blackwell and Good, and the epistemic version due to Wayne Myrvold (2012); in the second half of the paper, I put this framework to work, turning to some of the questions from the recent debate about inquiry and asking how the Blackwell-Good-Myrvold approach can help us answer them. Questions will include: Are there purely epistemic norms that govern these actions (Flores and Woodard forthcoming)? When should we initiate an inquiry, when should we continue it, when should we conclude it, and when should we reopen it? How should we understand Julia Staffel's distinction between transitional attitudes and terminal attitudes (Staffel 2021a,b)? How do epistemic norms of inquiry relate to epistemic

*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; and 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 is an extended version of the talk I gave there.

norms of belief or credence, and can they conflict (Friedman 2020)?
And how should we understand the epistemic error that occurs when
someone is resistant to evidence (Simion 2023)?

In the recent philosophical literature on inquiry, epistemologists point out that their subject has often begun at the point at which you already have your evidence, and it has then focussed on identifying the beliefs for which that evidence provides justification or which count as knowledge for someone with that evidence.¹ However, this recent literature goes on to point out, we are not mere passive recipients of the evidence we have. 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; we'll prod the world in one way to see how it responds but we won't prod it in another; 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 we 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 normativity of inquiry.

In this paper, I approach these questions from the so-called formal side of epistemology, which is to say the side that has traditionally talked of credences rather than beliefs. Developing out of the philosophy of science, as it did, this branch of epistemology has long discussed inquiry. And, building on the insights of the statistician David Blackwell (1951) and the mathematician I. J. Good (1967), it has produced a reasonably well-developed framework in which to understand norms of inquiry, both epistemic and practical. In the first half of the paper, I will present this framework and the pragmatic version of the so-called Value of Information Theorem that is associated with I. J. Good (Section 1); I will describe some generalizations of that result, and I will explain its limitations. Then I will present an epistemic version of it due to Wayne Myrvold (2012), generalize that in similar ways, and again look at limitations (Section 2). In the second half of the paper, I put this approach to work, turning to some of the questions from the recent debate about inquiry and asking how the Blackwell-

¹A small sample of recent writings: (Hookway, 2006; Friedman, 2020; Kelp, 2021; Thorstad, 2022; Simion, 2023; Flores & Woodard, forthcoming).

Good approach can help us answer them. Questions will include: Are there purely epistemic norms that govern these actions, as Carolina Flores and Elise Woodard (forthcoming) contend (Section 3)? When should we initiate an inquiry, when should we continue it, when should we conclude it, and when should we reopen it (Section 4)? How should we understand Julia Staffel’s distinction between transitional attitudes and terminal attitudes (Staffel, 2021b) (Section 5)? How do epistemic 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 (Section 6)? And how should we understand the epistemic error that occurs when someone is resistant to evidence in the way Mona Simion (2023) describes (Section 7)?

Throughout, I will present the ideas both informally and formally. I’ll restrict the formal presentation using mathematical notation to the grey boxes. These can be skipped over, if you prefer, as the ideas are presented informally in the surrounding text.

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—see, for instance, (Zollman, 2007, 2010; Rosenstock et al., 2017). 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.

1 The pragmatic value of gathering evidence

In the framework that Blackwell and Good introduce, we represent an individual’s doxastic state not by the beliefs they have, but by their credences or partial beliefs, as is reasonably standard in formal epistemology. These are the states we report when we say ‘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. We collect together all the propositions to which an individual assigns a credence and call it their *agenda*. For ease of exposition, we assume there is a finite set of possible worlds such that each proposition in the agenda can be represented by the set of these worlds at which it’s true. Our individual’s *credence function* is then the function that takes each proposition in the agenda and assigns to it the number at least 0 and at most 1 that represents their credence in it—0

represents minimal credence; 1 represents maximal credence. We'll assume throughout that our inquiring individual's credence function at any time is probabilistic. That is, it assigns 1 to all necessary truths, 0 to all necessary falsehoods, and the credence it assigns to a disjunction of two mutually exclusive propositions is 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 $C : \mathcal{F} \rightarrow [0, 1]$.

And a credence function is *probabilistic* iff

- (i) $C(\emptyset) = 0$ and $C(W) = 1$;
- (ii) $C(X \cup Y) = C(X) + C(Y)$ if $X \cap Y = \emptyset$.

Equivalently,

- (i) $\sum_{w \in W} C(w) = 1$;
- (ii) $C(X) = \sum_{w \in X} C(w)$.

1.1 The framework

While most discussion of Good's 1967 paper, 'On the Principle of Total Evidence', focuses on what has become known as the Value of Information Theorem, the real contribution lies in his definition of the pragmatic value of gathering evidence, which was anticipated by David Blackwell in his 1951 paper, 'Comparison of Experiments' (Blackwell, 1951; Good, 1967).

This definition begins with another definition; it begins with a definition of the pragmatic value of a credence function. 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 it by the credence you assign to that world, and summing up these credence-weighted utilities; and then you pick an option whose expected utility is maximal. 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 *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 τ , is*

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

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 .

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.

Now, having defined the pragmatic value of a credence function relative to a particular decision you'll face and a way of breaking ties, Good 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* $\mathcal{E} : \mathcal{W} \rightarrow \mathcal{F}$. This takes each world w to the 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 , your posterior at world w will be $C(- \mid \mathcal{E}_w)$, providing $C(\mathcal{E}_w) > 0$.

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*

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

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.

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 Good's 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

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

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

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

1.2 The theorem

Good’s Value of Information Theorem then runs as follows: Fix a decision problem you’ll face at a later time; fix the way you 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 evidence that will teach you which element of a particular partition is true—perhaps you’ll learn that the forecast says rain or the forecast says dry and those are the only possibilities. 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, if you assign some positive credence to a state of the world in which the evidence you’ll learn will change how you make the decision you’ll face, 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 change your mind.

Good'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 $\tau(\mathcal{D}_{C(-|\mathcal{E}_w)}) \neq \tau(\mathcal{D}_C)$.

A natural response upon first encountering Good's Theorem is to think it's obviously true. After all, surely any true evidence is guaranteed to improve our epistemic situation, and surely improving the epistemic standpoint from which you face decisions leads to better choices. But 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. 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 than choosing.

So Good's Theorem is not obviously true. It holds because, 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 generalize Good's result in the next section.

1.3 Generalizing the theorem

Now, Good's Theorem 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 Good's theorem is interesting, the real value of the framework that Blackwell and he introduced is that it allows us to assign a pragmatic utility to an evidence-gathering episode. 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. 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.

What's more, it might be that some of the alternative options available to me aren't evidence-gathering episodes at all. But the account of pragmatic value of evidence-gathering that we obtain from Blackwell and Good allows us to compare them anyway. 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, Blackwell and Good assume you know for sure which decision you'll face using your credences, but of course you might be uncertain of this. But again, it's easy to incorporate this: we simply ensure 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 ensure that we assign credences to these more fine-grained possible worlds; and, having done all this, we can define the pragmatic value of a credence function at a 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; and then the pragmatic value of an evidence-gathering episode is again the pragmatic value of the credence function you'll end up with after gathering the evidence and updating on it. And Good's theorem still goes through with this amendment.

#4: Allowing non-factive, non-partitional evidence. As stated, Good's 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 true 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 Good's theorem still hold? Can we find weaker conditions on evidence-gathering episodes such that Good's theorem still holds?

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 Good's 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. 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.)

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

Theorem 2 ((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 \mathcal{W} such that $C(w) > 0$ and $f(\mathcal{D}_{C(-|\mathcal{E}_w)}) \neq f(\mathcal{D}_C)$.

Kevin Dorst et al. (2021) approach the generalization in a slightly different way. Where Geanakoplos places a condition directly on the evidence-gathering episode, Dorst et al. place conditions on the prior credence function and its attitudes to the possible posterior credence functions you would have were you to learn and update on the evidence. They show that, if it satisfies those conditions, then for any decision problem you'll face, gathering the evidence is never worse and often better than not gathering it, in expectation.

Dorst et al.'s strengthening of Good's Theorem

Some preliminary definitions:

- 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 .
- Given an updating plan R , a random variable X , and a real number t , let $\langle \text{Exp}_R(X) \geq t \rangle$ be the proposition that is true at all worlds w for which $\text{Exp}_{R_w}(X) = \sum_{w' \in W} R_w(w')X(w') \geq t$.
- Given a credence function C and an updating plan R , we say that C *totally trusts* R if, for any random variable X and any threshold t , the following holds:

$$\text{Exp}_C(X \mid \langle \text{Exp}_R(X) \geq t \rangle) \geq t$$

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

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

Theorem 3 (Dorst et al., 2021). *If C totally trusts R , 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 \mathcal{W} such that $C(w) > 0$ and $f(\mathcal{D}_{R_w}) \neq f(\mathcal{D}_C)$.

I won't delve 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 why 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.

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. It is the sort of case that Nilanjan Das (2023) calls 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.

Should you take this evidence? Well, it very much depends on your prior and the decision you'll face with your posteriors. 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. And indeed Das shows that, for biased inquiries, there is always some prior and decision problem that will lead you to rationally reject such biased inquiries. 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.

As 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 op-

tion 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.

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.

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 out our evidence being different from how it actually is. A neat illustration is Tim Williamson's 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 only has a sweeping single hand. You know it's either 12noon, 1pm, . . . , 10pm, or 11pm, but you know don't know which. If it's 1pm, your evidence will be that it's 12noon, 1pm, or 2pm; if it's 2pm, your evidence will be that it's 1pm, 2pm, or 3pm; and so on.

Should you look at the clock? Again, it very much depends on your priors and the decision you'll face. Suppose you currently assign equal credence to each possible time. The first option available pays a million dollars if it's 12noon, 1pm, . . . , 5pm, and nothing otherwise; the second pays a million dollars if it's 6pm, 7pm, . . . , 11pm, and nothing otherwise. Then you should gather the evidence. But, as Nilanjan Das (2023) notes, if the first option pays a million dollars if the time is an even number and the second pays a million dollars if the time is 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.

#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 Good's framework to argue for this Bayesian assumption, at least in those cases to which Good's 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 it's 2pm and I receive the evidence it's either 1pm, 2pm, or 3pm, I should conditionalize not on this evidence, but on the fact I received it, which is true only at 2pm.

There is an interesting ongoing debate whether such an updating plan is really available to me. You might think it is not, since it requires me to reflect on the evidence I in fact have, and then infer the worlds at which I would receive that. But of course I'm not supposed 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 the evidence is. I won't delve deeper into this debate here (Carr, 2021; Gallow, 2021; Isaacs & Russell, 2023; Schultheis, ta).

However, it is worth noting that, if Schoenfield's updating rule is genuinely available to us, then Good's 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 $\mathcal{E}_w = \mathcal{E}_{w'}$, $R_w^\mathcal{E} = R_{w'}^\mathcal{E}$.
- 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 4 ((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 such that*

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

- (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 $R_w(-) = C(- \mid \mathcal{E}_w)$.

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 how to choose.

2 The epistemic value of gathering evidence

Good’s theorem tells us something about when you have *practical* reason to engage in a certain sort of evidence-gathering. But, as Wayne Myrvold (2012) shows, building on work by Graham Oddie (1997) and Hilary Greaves and David Wallace (2006), there is also a version that tells us something about when you have *epistemic* reason to gather evidence. Alejandro Pérez Carballo (2018) has extended Myrvold’s approach in various ways.

Recall: Good’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 says 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 Richard Pettigrew (2023) have recently defended it in a different way.

Strictly proper epistemic utility functions

An *epistemic utility function* EU takes a credence function C and a possible world w and returns $EU(C, w)$, 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 all probabilistic credence functions P and alternative credence functions $C \neq P$,

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

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 cre-

dence 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, 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 $\text{Brier}(C, w)$ 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.

Definition 5 (Enhanced log score). *The enhanced log score $\text{Log}(C, w)$ 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.

So now we have a way of assigning epistemic value to a credence function at a world. And so we can simply appeal to Good's insight to 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 Myrvold's epistemic version of Good's 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.

Myrvold's Epistemic Value of Information Theorem

Theorem 5 ((Myrvold, 2012)). *If EU is strictly proper and \mathcal{E} is factive and partitional,*

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

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 to propositions that are true at that world and minimal credence 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 Good's result, the reason that Myrvold'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. As with Good's result, Myrvold's 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 would happen, is outweighed by the epistemic improvements that are possible, once those are weighted by the prior's probability that they would happen instead.

As with Good's 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 Myrvold's Theorem

Theorem 6 ((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 ,*

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

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

Dorst et al.'s strengthening of Myrvold's Theorem

Some preliminary definitions:

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

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

- 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 $R_w(X) \geq t$.
- Given a credence function C and an updating plan R , we say that C *simply trusts* R if, for any proposition X and any threshold t , the following holds:

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

Theorem 7 ((Levinstein, 2023)). *If C simply trusts R , then for any additive and continuous strictly proper epistemic utility function EU,*

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

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.

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. Interestingly, relative to the enhanced log score, this isn't the case. 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. And this time, the same is true for the enhanced log score.

Finally, we conclude this tour of the landscape that has sprung up around Good's theorem by pointing to the epistemic analogue of the argument for Miriam Schoenfield's version of conditionalization. Regardless of the evidence-gathering episode you face, 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 (Greaves & Wallace, 2006; Schoenfield, 2017).

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

Theorem 8 ((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 ,*

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

- (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 ,*

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

3 Are there epistemic norms of evidence-gathering?

In the second half of this paper, I turn from the formal results and arguments that extend Good's theorem to the recent literature on inquiry that has been developing in mainstream epistemology. My plan is to apply the insights from the first part to answer some of the central questions that have arisen.

The first questions: Are there epistemic norms that govern evidence-gathering? Are there epistemic reasons to gather evidence?

Carolina Flores and Elise Woodard argue that there are such epistemic norms. They argue that this is the best explanation of the fact that we engage in a distinctively epistemic sort of criticism when someone fails to inquire properly. We do this, they think, when someone doesn't draw their evidence from a sufficiently diverse range of sources, when they are resistant to evidence that is easily available, when they are simply too lazy to

gather evidence that is at hand, and so on.

I agree with them that we engage in such criticism. But I don't think those of us who agree with their conclusion can appeal to this fact to argue in favour of it. After all, those who deny that there are epistemic norms of inquiry will not deny that we criticize people who inquire poorly; they'll just deny that there's anything distinctively epistemic about those criticisms; or they'll deny that whatever is distinctively epistemic about them gives us reason to think we are criticizing them for breaking epistemic norms or failing to respond to epistemic reasons. They will say that, just because the criticism talks of standardly epistemic phenomena, such as evidence and beliefs, that does not mean it is based on a breach of an epistemic norm in the way that our criticism is based on such a breach when we say that someone has not proportioned their beliefs to the evidence.

Nomy Arpaly (2023) is someone who might respond in this way. She argues that there are no epistemic norms of inquiry because there are no epistemic reasons for inquiring. What there are instead are *instrumental* reasons for inquiring *for someone whose goal is knowledge or accuracy or some other putative aim for belief*. If your goal is knowledge, you should gather a diverse range of evidence; if your goal is accurate belief, you should not leave free evidence on the table; and so on. There are norms like this, Arpaly thinks. But these are not epistemic norms any more than the following is an aesthetic norm: If you want a lovely house, you should save money until you can afford one. They are instrumental norms that explain what means we have reason to take if we have certain ends.

I think Myrvold's version of Good's theorem, and the various generalizations I described above, suggest a norm for inquiry that runs as follows:

The Epistemic Norm of Inquiry Gather evidence so as to maximize the expected epistemic utility of your future credence function.

And I think this is an epistemic norm for inquiry. But Arpaly will argue that it is really an instrumental norm directed at those who value accuracy. My response is that it is both, and indeed that, in the end, all epistemic norms are instrumental norms for those who value accuracy. This is the core tenet of accuracy-first epistemology with its veritist axiology for beliefs and credences and its teleological conception of rationality and epistemic normativity.

We can get a sense of what Arpaly takes an epistemic reason to be by considering her Sinking Heart Intuition, which she uses to argue against Susanna Rinard's (2017) claim that there are practical reasons for believing.

Imagine that you have cancer and you do not yet know if the course of chemotherapy you have undergone will save you or not. You sit down at your doctor's desk, trying to brace yourself

for news, aware that at this point there might be only interim news—indications that a good or a bad outcome is likely. The doctor says there are reasons to be optimistic—to believe that everything will come out OK. Though you are still very tense, you perk up and you feel warm and light all over. You ask what the reasons are. You’re all ears. In response, the doctor tells you about ironclad scientific results showing that optimism is good for the health of cancer patients.

Your heart sinks. You experience a very bitter disappointment and will probably be angry at the doctor for the misleading way he put his point. (Arpaly, 2023)

What the doctor gives you is not an epistemic reason for optimism—that is, for high credence that the chemo worked. What he gives you is a practical reason for trying to bring yourself to be optimistic. For Arpaly, an epistemic reason for optimism, in contrast, is the sort of thing you were expecting to hear when he said that, namely, some fact that stands in the evidential support relation to the proposition that says the chemo worked.

So Arpaly would like to distinguish a reason to believe a proposition, which is simply a fact that supports that proposition, from a reason to do things that will lead you to have a belief with certain desirable properties, such as being accurate or counting as knowledge. But this distinction turns on the assumption that there is some objective relation of evidential support that holds between propositions: *It seems that there is an external world* supports *There is an external world*; *Ada said that it is raining* supports *It is raining*; *The CT scan shows no tumour* supports *The chemo worked*; and so on. But this, it seems to me, is an illusion. In the end, there is a range of rationally permissible prior credence functions—how wide depends on just how radical your subjectivism is; but there is more than one. One proposition supports another relative to a prior credence function if the credence assigned to the second conditional on the first is greater than the unconditional credence assigned to the second—and this is the only sort of evidential support relation there is; the only objective evidential support facts are those that hold relative to any rationally permissible prior, such as that a conjunction supports each of its conjuncts, or perhaps that *The objective chance of rain is high* supports *It will rain*. What determines which prior credence functions are rationally permissible are considerations of accuracy.² And so, ultimately, to the extent there are facts about evidential support and therefore epistemic reasons of the sort that Arpaly seeks and so epistemic norms in her sense, they are ultimately just instrumental reasons and instrumental norms that govern those who value accuracy. And so they are of the same type as The Epistemic Norm of Inquiry that I stated above.

²See (Pettigrew, 2022) for a recent attempt to establish this sort of epistemic permissivism about credences from an accuracy-centred point of view.

Of course, all of this is contentious. It requires a much longer argument to establish any of this. But, regardless of the merits of that argument, I'd want to hear a lot more about what could possibly ground the objective evidential support relation posited by those who build a notion of epistemic reason on it before being convinced that there is a real distinction here.

On the view I'm sketching, what makes a reason epistemic is that it is grounded in the epistemic value of doxastic states that are closely connected to whatever the reason is a reason for; and what makes a norm epistemic is that it holds because of facts about epistemic value. For me, those facts are facts about the accuracy of the belief; but for others they might be something else, such as knowledge or understanding or wisdom.

One worry about this sort of view, which Sophie Horowitz has raised and that Flores and Woodard discuss, is that it over-generates epistemic reasons and epistemic norms. A detective is settling in for a long night working through the evidence against a suspect in order to decide whether or not to charge them in the morning. If she keeps drinking coffee, she'll power on through to the early hours and read all of the relevant evidence; if she doesn't, she'll fall asleep at her desk at 4am and miss out on much of the evidence. On the view that Flores and Woodard favour, and with which I agree, the detective has epistemic reason to take in all of the evidence she can. But if that's so, Horowitz challenges, surely she also has epistemic reason to drink coffee. And surely that's absurd! While she has *instrumental* reason to drink the coffee, given that she values being accurate and drinking coffee will serve that end, she doesn't have *epistemic* reason to do so.

In the end, I'm happy to bite the bullet here. The notion of an epistemic reason is not really a pre-theoretic one; it is closer to a technical notion used in epistemology. So I don't think we need to try too hard to respect whatever our philosophical intuitions are concerning its usage. And in the end, what harm is done by allowing lots of reasons to count as epistemic? It is not an accolade whose prestige we must preserve by awarding it only sparingly.

A nice feature of The Epistemic Norm of Inquiry is that it can account for the epistemic criticisms we level at the different characters that Flores and Woodard describe. Let's consider just two:

Gullible Gabe Gabe tells you that there are 10% fewer jobs in finance this year than there were last year. You defer to him. You later learn that he got this fact from a dated *Economist* magazine that he read at his therapist's office, assuming that it was up-to-date despite the prevalence of dated magazines in therapists' office—something he should know about.

In fact, I don't think we have quite enough information from this vignette to direct our criticism of Gabe precisely. It might not be the way he gath-

ered evidence that is ultimately at fault, but rather his high credence that the source from which he gathered evidence is reliable. From the point of view of a high credence that the source is reliable, the expected pragmatic or epistemic value of gathering evidence from it is likely to be high, and therefore rationally required, at least from the point of view of the agent's own credences, given there are few costs. But that original high credence itself might be irrational because it isn't a good response to Gabe's evidence concerning the source's reliability; or it might be a good response to the evidence Gabe in fact has about the source's reliability, but it is nonetheless flawed because he should have gathered further evidence about this in the past and failed to do so.

This reveals an interesting set of distinctions. When we use the Good-Myrvold framework, there are different perspectives from which we might assess the epistemic rationality of Gabe's evidence-gathering behaviour. Consider two of the key components in this framework: first, the credences from the point of view of which we assess Gabe's evidence-gathering; second, the pragmatic utilities that specify the cost to Gabe of gathering that evidence. Recall that, even in Myrvold's version of the framework, where we assess evidence-gathering episodes for their purely epistemic value, we must still specify the cost in terms of pragmatic value and give an exchange rate between the two sorts of value in order to give an all-things-considered assessment of whether or not to gather certain evidence. Now, while I presented Good's and Myrvold's frameworks as if we always use the evidence-gatherer's *actual* prior credences and the evidence-gatherer's *actual* pragmatic utilities to assess these episodes, we can easily substitute in different priors and different utilities. So, for instance, we might use Gabe's actual prior credences, but we might also use the credences Gabe would have had if he had properly responded to the evidence that he's gathered in the past; or we might even use the credences Gabe would have had if he had properly responded not just to the evidence he's actually gathered in the past, but to the evidence he should have gathered in the past. And we can assess these evidence-gathering episodes not just from the point of view of his actual pragmatic utilities, but from the point of view of the utilities we think he should have.

How might these play out in this case? Well, perhaps relative to his actual high credence that magazines in therapist's offices are reliable, his evidence-gathering was rational, but that actual high credence is the result of him not gathering evidence in the past that, from the point of view of his expected epistemic utility at the time, was irrational. Then we might say either (i) his current evidence-gathering is rational, but it's based on credences that result from previous irrational evidence-gathering; or (ii) we might say his current evidence-gathering is irrational, because it doesn't maximize expected epistemic value from the point of view of the credences he would have had if he had gathered evidence rationally in the past. For

my money, I don't think it matters too much which of these we favour, as long as we're clear about the point of view from which the assessment takes place.

Let's turn now to another of Flores and Woodard's examples:

Lazy Larry Larry is a chemistry major, who forms his beliefs about the structure of the atomic nucleus based on over-simplifying and idealizing diagrams, depicting electrons as marble-like entities that orbit the nucleus in precise tracks. However, this is misleading: electrons actually are spread out diffusely within a massive region. The textbook includes this information, but Larry limits his efforts to just looking at the pictures.

In this case, it's important that Larry ignores the further evidence his source can provide as a result of laziness and not simply because he lacked the time. If he had only a little time with his source and couldn't attend to all the evidence it provides, then there might be nothing wrong with his evidence-gathering—he did the best he could under the constraints placed on him! But, as Flores and Woodard describe the case, he could have gathered more evidence, but he failed to do so. In that case, it's likely that the cost of gathering that extra evidence and updating on it is considerably less than the epistemic value he expects to get from it. And this explains why he is criticizable: he violates the Epistemic Norm of Inquiry.

Another possibility in Larry's case is that he simply values epistemic value rather little and values idle moments rather a lot! That is, the way he sets his exchange rate between epistemic value and pragmatic value, the pragmatic value he loses by attending to the information in the textbook when he could be idly staring into space is rather high, and considerably higher than the epistemic value he'd gain from the reading. In this case, it would be rational from his own subjective point of view not to gather the evidence. In that case, if we wish to judge him to be in error in some way, we might assess him not from the point of view of his own credences and utilities, but from the point of view of his own credences and the utilities we feel he should have. Perhaps we think it's a prudential or even moral failing to value idleness.

4 When should we start, continue, conclude, and re-open inquiry?

So far, we have just been talking about evidence-gathering episodes, and not inquiry. But an inquiry is simply a sequence of such episodes. We often embark upon an inquiry with the aim of answering some question, but equally we sometimes just inquire about a certain topic because we find

it interesting and are happy to learn what we learn without specifying a precise question to which we wish to find the answer.

About inquiry, we can ask: when should we embark on it? when should we continue an inquiry on which we've embarked? when should we conclude one? when should we re-open one? The Good-Myrvold framework gives us a neatly unified answer to all: from the pragmatic point of view, you should inquire if doing so maximizes expected pragmatic utility; from the epistemic point of view, you should inquire if doing so maximizes expected epistemic utility; and, once you've fixed an exchange rate between epistemic and pragmatic utility, you can define all-things-considered utility and say that you should inquire if doing so maximizes that quantity in expectation.

We will return to this point below, when we discuss Julia Staffel's notion of transitional and terminal attitudes, but it's worth noting here: 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 your credences in expectation, 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 them will make things worse. Provided you're confident enough that you'll see the series through to the end, and providing 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. We'll see an example below that seems to have that structure.

When should you cease inquiring further? From the pragmatic point of view, 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. This happens if you care only about the pragmatic value of your credences as a guide to action in the face of the decision you know you'll face with them. At some point, you 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 will 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. This can lead you to abandon before they're complete inquiries that it was nonetheless rational to embark on in the first place: this can happen because the costs of gathering further evidence in that inquiry has 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, the 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. Indeed, you will rarely acquire such certainty. 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 in the mainstream literature on inquiry by Christopher Willard-Kyle (forthcoming), who argues that, in inquiry, there will nearly always be room for improvement from an epistemic point of view. Willard-Kyle is 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. He argues 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.

5 Do we have transitional attitudes?

Recently, Julia Staffel (2019, 2021b,a) has drawn an interesting distinction between what she calls *transitional* and *terminal attitudes*. On her account, during the course of an inquiry, we form transitional versions of the attitudes we seek, whether these are outright beliefs or precise credences. Only when the inquiry is complete do we form terminal versions of those attitudes. So, for instance, a detective who is methodically working her way through the body of evidence her team has amassed forms transitional credences concerning the identity of the culprit, and only after she has surveyed all this evidence does she form terminal credences on that matter. Staffel says that what distinguishes these attitudes is what we're prepared to do with them: we are prepared to act on terminal attitudes but not on transitional ones; we are prepared to make assertions based on terminal attitudes but not on transitional ones; and we are prepared to feed terminal attitudes into future deliberation and reasoning, but we are not prepared to do the same with transitional attitudes. In the section, I want to explore how Good's and Myrvold's results can shed light on this distinction.

Let me begin in a surprising way by offering an argument that there can be no transitional attitudes that answer to Staffel's description. Let's suppose I face a decision in the midst of my inquiry that I expected to face at the end. In this case, it seems that I have no choice but to choose using the credences I have at that point, which have been obtained from my credences at the beginning of the inquiry by updating on the evidence I've received during its course to date. After all, what else is available to me? Of course, there are my credences at the beginning of my inquiry. Should I use those instead? The problem with that suggestion is that those credences themselves don't think I should use them, at least if the evidence-gathering episodes I've embarked on so far are ones that the prior credences expect to have greater pragmatic value than not embarking on them, such as if those evidence-gathering episodes have the features that make Good's theorem or Geanakoplos' generalization of it applicable. Sure, my prior credences would have liked it even more if I'd got to complete my inquiry before facing this decision, but the world has prevented that and I must act now. So, if I must act either on my priors or on my current credences, which I hold mid-inquiry, I should act on my current ones, which suggests they're not transitional. Indeed, if I choose using my prior credences, I will violate the Principle of Total Evidence. After all, while I might not have gathered all the evidence I wanted to gather before I had to face the decision, I did gather some, and choosing using my priors from the beginning of the inquiry is to ignore that evidence.

What's more, the same goes for making assertions on the basis of those credences and feeding them in to further deliberation and reasoning: even if my inquiry is interrupted before it's complete, I should nonetheless choose

whether and what to assert based on the credences I have reached by that point, and it is those credences I should feed into future deliberation and reasoning. To make an assertion is simply to undertake an action of a certain sort and its correctness can be evaluated in the same way as other actions by looking at its expected utility; so the analysis just given applies as much to the decision whether and what to assert as to other decisions. And surely when we feed attitudes into our further deliberation and reasoning processes, we want to satisfy the Principle of Total Evidence, and so it is again the credences we've reached mid-inquiry, before that inquiry was interrupted, that we should use.

Now, this argument doesn't quite establish that there are no transitional attitudes of Staffel's sort. If an inquiry is made up of a series of evidence-gathering episodes, each of which satisfies the conditions that make Good's theorem (or Geanakoplos' generalization) hold, then the argument suggests that no attitudes formed during that inquiry have the hallmark features of the transitional. But, as we saw above, not all inquiries are like that. Sometimes the whole sequence of evidence-gathering episodes is such that we expect our credence function to be better after they're all completed, but there are points in the course of the investigation at which we expect our credence function will be worse. This might happen, for instance, if we string together a bunch of Das' biased inquiries, where those in the first stretch are biased in one direction and those in the second are biased in the other, but taken together, they aren't biased in either direction.

For instance, take an example that Staffel considers in her recent book manuscript. A detective asks her team to divide up the evidence they've gathered into that which suggests the first suspect is guilty and that which suggests the second culprit is guilty. She plans to work through the first set first and the second set second. In this case, while her prior credences expect the credences she'll have once she's worked through both sets to be better than they expect themselves to be, they also expect the credences she'll have once she's only worked through the first set to be worse than they expect themselves to be. And so, if she's interrupted just as she completes the first set and suddenly has to make a decision she was hoping to make only at the end, she might well decide not to use her current credences. And in that sense they are transitional.

So, while I think the argument against transitional attitudes that I gave above fails to show there are no such attitudes, I think it might show that they're rather rarer than Staffel imagines. Most planned inquiries involve a sequence of evidence-gathering episodes each of which improves your credences in expectation; the minority are like the detective running through a series of individually biased but collectively unbiased inquiries. Note that she had to ask for the evidence to be curated in a particular way to even achieve that sort of biased inquiry. And also note that it's rather difficult to fill in the details of the detective's case in such a way that it is composed

of biased inquiries. After all, she proceeds by looking at different pieces of evidence in turn. Her team has curated these for her and determined the order in which she'll look at these different pieces of evidence, but the detective knows this and knows what the possibilities are in advance. So, for instance, when she picks up the file containing the first piece of evidence, she knows what she might find in there. Let's say there are just four options: DNA evidence that implicates the first suspect, a confession by the first suspect, a discrepancy in the statement of the first suspect, or a witness statement identifying the first suspect. And, what's more, the detective's evidence is luminous to her—this is not a case like Williamson's unmarked clock or the good case vs bad case. So whichever she learns, she also learns that this is what she's learned. But in that case, the evidence-gathering episode is factive and partitional. And so it satisfies the conditions required for Good's theorem. What's more, due to the law of total probability, it cannot be that, whichever piece of evidence she receives will raise her credence in the guilt of the first suspect.³ And in any case, since it's factive and partitional, she expects that gathering the evidence improves her epistemic situation. Of course, this is just one example. But I hope it illustrates that it's actually reasonably difficult to come up with inquiries on which we intentionally embark that are biased in the way required to violate Good's theorem and give rise to genuinely transitional credences.

Staffel wishes to apply her account of transitional and terminal attitudes not only to cases of empirical inquiry, like the detective investigating a murder, but also to cases of logical reasoning and other *a priori* inquiry, like the logic student who uses truth tables to establish whether $(p \rightarrow (p \rightarrow p))$ is a tautology or not. Following Ian Hacking (1967), as well as recent developments of his view by Robbie Williams (2018) and Richard Pettigrew (2020), we might appeal to Good's approach to inquiry to model logical and *a priori* reasoning as well.

The idea is that, just as empirical evidence serves to rule out certain possible worlds as not actual, so the fruits of logical inquiry and other forms of *a priori* reasoning also serve to rule out worlds as not actual. Now, the worlds that they rule out are not possible worlds in the standard sense. When the logic student completes the first row of the truth table for $(p \rightarrow (p \rightarrow p))$ and thereby discovers that this formula is true when p is true, this rules out the world in which that formula is false when p is true. But this isn't a genuine possible world, because it isn't logically possible. Rather, Hacking suggests, it is a *personally possible world*: that is, it is possible from the point of view of the student's epistemic position. According

³The law of total probability says that her credence in the guilt of the first suspect should be her expectation of her conditional credence in their guilt given each of the possible pieces of evidence. If her conditional credence in their guilt is greater given each possible piece of evidence, the expectation is greater than the original credence, and that violates the law of total probability.

to Hacking, we might represent the student as distributing credences over these personally possible worlds in such a way that they sum to 1 and her credence in any proposition is the sum of her credences in the personally possible worlds at which that proposition is true. In this way, we can give a sort of Bayesian representation of a reasoner who is not logically omniscient. And then we can understand logical learning in the same way the Bayesian understands empirical learning: both rule out worlds; and, having ruled out worlds, we assign zero credence to them and then scale up our credences in the remaining worlds so that they again sum to 1. On this picture, logical learning is formally represented exactly as empirical learning is, and so Good's theorem and Geanakoplos' generalization, as well as Myrvold's theorem and its generalizations, all hold of logical learning just as they do of empirical learning. We might represent our logic student as having a prior credence that $(p \rightarrow (p \rightarrow p))$ is a tautology; then, when she learns the formula is true when p is true, she rules out worlds at which it is false when p is true and updates her credences; and then, when she learns it's true when p is false, she updates her credences again, becoming certain, or near to certain, that it's a tautology.

If this is the right way to represent logical ignorance and logical learning, then the same points apply to transitional attitudes within logical or other *a priori* reasoning that I made about such attitudes in the midst of empirical inquiry above.

6 Is there a tension between the epistemic and the zetetic?

Let's turn now to an example that motivates much of Jane Friedman's recent contributions to the literature. I'll quote at length:

I want to know how many windows the Chrysler Building in Manhattan has (say I'm in the window business). I decide that the best way to figure this out is to head down there myself and do a count. [...] Say it takes me an hour of focused work to get the count done and figure out how many windows that building has. [...] Now think about the hour during which I'm doing my counting. During that hour there are many other ways I could make epistemic gains. [...] First, I'm a typical epistemic subject and so I arrive at Grand Central with an extensive store of evidence: the body of total evidence, relevant to all sorts of topics and subject matters, that I've acquired over my lifetime. Second, I'm standing outside Grand Central Station for that hour and so the amount of perceptual information available to me is absolutely vast. [...] However, during my hour examining the Chrysler Building I barely do any of that. I need to get my count right, and to do that I really have to stay focused on

the task. Given this, during that hour I don't extend my current stores of knowledge by drawing inferences that aren't relevant to my counting task, and I do my best to ignore everything else going on around me. And this seems to be exactly what I should be doing during that hour if I want to actually succeed in the inquiry I'm engaged in. [...] There is an important sense in which I succeed in inquiry by failing to respect my evidence for some stretch of time. It's not that my success in this case comes by believing things my evidence doesn't support, but it does come by ignoring a lot of my evidence and failing to come to know a great deal of what I'm in a position to know. (Friedman, 2020)

I think the natural thing to say here is that, as Friedman faces the Chrysler Building, she faces a choice between a number of different evidence-gathering episodes she might undertake. Some of them are the ones that form the inquiry she is there to undertake, namely, determining the number of windows in the building; some involve attending to sensory information and perhaps testimony that is available at the spot where she's ended up, but which is irrelevant to her inquiry; and some involve drawing inferences from the store of memories and other evidence she's previously collected, which is again irrelevant to her inquiry.

Of course, it's rather unusual to think of these last episodes as involving evidence-gathering. After all, you already have the evidence, and you're simply drawing conclusions from it that you haven't drawn before. But, as I described in the previous section, I think it's reasonable to view logical reasoning as doing something similar to what gathering empirical evidence does. In both cases, they are ruling out states of the world that are in some sense possible.

So, having seen this, we can understand the logical reasoning that Friedman doesn't do when she's in front of the Chrysler Building as just another sort of evidence she doesn't gather, just as she doesn't gather the evidence she might do if she were to attend to the conversation between the two commuters standing to her left, say. And once we do that, we can say that Friedman does the right thing by continuing with her window-counting inquiry so long as, at each stage, the evidence-gathering episode that comes next in that inquiry is the one that maximizes expected pragmatic or epistemic value among those episodes that are available to her. And if we see things in this way, there is no clash between an epistemic norm and a zetetic one. There are just one norms: gather evidence in the way that maximizes expected utility; and respond to the evidence you gather by conditionalizing. And they govern what Friedman should do in front of the Chrysler Building.

It might seem that this all rather misses the point. In Friedman's example, you might think, the question is not whether to gather certain evidence

that surrounds you as you stand outside Grand Central Station. The point is that you already have that evidence, whether you want it or not, simply by being there. You cannot help but have it. And so the real question is whether to incorporate it or not when doing so might take up resources that can be used to gathering the evidence about the number of windows. And that's what creates the clash between epistemic and zetetic norms, since Friedman thinks you should not incorporate the evidence because of how doing so interferes with your inquiry, but presumably standard epistemic norms say you should incorporate that evidence, since you have it.

In the end, this comes down to when you want to say that someone has certain evidence, and what exactly that means. The fact that all this evidence is to hand around me outside Grand Central Station does not mean that I have it. It just means that it's easily accessible to me, should I wish to gather it. But even if we wish to say that we have some of it because, perhaps, it's impossible to ignore certain things like the content of a very loud conversation or the presence of a very brightly coloured car nearby, we can model this in a way that allows us still to use Good's framework. We can say that the evidence is now stored in us somewhere and somehow, but we haven't yet brought it to our attention; we haven't yet passed it to whatever part of our mind takes in evidence and alters credences in the light of it. Rather, it's sitting in storage waiting for us to decide whether or not to attend to it. But then the decision whether or not to attend to it is just like the decision whether or not to gather and update on some evidence, and Good's theorem applies. So again, I can know that there is the content of some loud conversation sitting in storage, and I can choose to bring it out of storage and attend to it or continue with my enumeration of the windows in the Chrysler building. And Good's theorem will tell me what to do.

7 Why should we not resist evidence?

In this final section, I want to ask what Good's 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 Good's 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 Good's 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—as I mentioned in my discussion of Nomy Arpaly's position above, I'm pretty sceptical of such evidential probabilities. 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.

So, for the subjective Bayesian, George is certainly flawed, but it's not because he is resistant to the evidence Anna gives him 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 it 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, and 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 different 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, leaving him with the biased credence function he has when he meets Anna and hears her testimony. In that latter case, he showed genuine resistance to evidence he received, and Myrvold's version of Good's 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, leading him to anguish, disorientation, and alienating him from the epistemic bubble into which this belief has drawn him.

Simion is right, of course, that Good's theorem seems to get this case wrong. 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, Good's theorem says, he does nothing wrong by not gathering it. And that, Simion contends, is the wrong answer.

I agree it is the wrong answer, but I don't think it's the one required by Good's theorem. As I mentioned above, while we typically apply Good's theorem using the individual's actual credences and their actual utilities, we needn't do that; we can instead apply it using their actual credences and the utilities *we think they ought to have*, for instance. So, in Jon's case, we can apply it using his credences and utilities that don't include this great cost to losing his belief; a cost that we feel shouldn't be factored in to the decision. And once we do that, Good's theorem delivers the result we want: Jon should gather the evidence. And so we can identify the flaw in his epistemic life.

8 Conclusion

The mainstream epistemology literature on inquiry is large and growing fast. I have only been able to consider a handful of the contributions in this paper. But I hope you'll find that the perspective offered by Good's theorem, Geanakoplos' generalization, and Myrvold's epistemic version and its generalization is illuminating for those questions. And this suggests, I think, that it will be similarly illuminating when applied to other questions from this literature.

References

- Arpaly, Nomy (2023). Practical reasons to believe, epistemic reasons to act, and the baffled action theorist. *Philosophical Issues*.
- Blackwell, David (1951). Comparison of Experiments. *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, 2, 93–102.
- Brown, Peter M (1976). Conditionalization and expected utility. *Philosophy of Science*, 43(3), 415–419.
- Carr, Jennifer (2021). A Modesty Proposal. *Synthese*, 198(4), 3581–3601.
- Das, Nilanjan (2023). The Value of Biased Information. *The British Journal for the Philosophy of Science*, 74(1), 25–55.
- Dorst, Kevin (2020). Evidence: A Guide for Uncertain. *Philosophy and Phenomenological Research*, 100(3), 489–757.
- Dorst, Kevin, Levinstein, Benjamin A., Salow, Bernhard, Husic, Brooke E., & Fitelson, Branden (2021). Deference Done Better. *Philosophical Perspectives*, 35, 99–150.

- Flores, Carolina, & Woodard, Elise (forthcoming). Epistemic Norms on Evidence-Gathering. *Philosophical Studies*, (pp. 1–25).
- Friedman, Jane (2020). The Epistemic and the Zetetic. *Philosophical Review*, 129(4), 501–536.
- Gallow, J. Dmitri (2021). Updating for Externalists. *Noûs*, 55(3), 487–516.
- Geanakoplos, John (1989). Game Theory without Partitions, and Applications to Speculation and Consensus. Tech. rep., Cowles Foundation Discussion Paper.
- Good, I. J. (1967). On the Principle of Total Evidence. *The British Journal for the Philosophy of Science*, 17, 319–322.
- Greaves, Hilary, & Wallace, David (2006). Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility. *Mind*, 115(459), 607–632.
- Hacking, Ian (1967). Slightly more realistic personal probability. *Philosophy of Science*, 34(4), 311–325.
- Hookway, Christopher (2006). Epistemology and Inquiry: The Primacy of Practice. In S. Hetherington (Ed.) *Epistemology Futures*, (pp. 95–110). Oxford University Press.
- Isaacs, Yoaav, & Russell, Jeffrey Sanford (2023). Updating without Evidence. *Noûs*, 57(3), 576–599.
- Joyce, James 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, Christopher (2021). *Inquiry, Knowledge, and Understanding*. Oxford, UK: Oxford University Press.
- Levinstein, Benjamin A. (2023). Accuracy, Deference, and Chance. *Philosophical Review*, 132(1), 43–87.
- Myrvold, Wayne C. (2012). Epistemic values and the value of learning. *Synthese*, 187(2), 547–568.
- Oddie, Graham (1997). Conditionalization, Cogency, and Cognitive Value. *British Journal for the Philosophy of Science*, 48, 533–41.
- Pérez Carballo, Alejandro (2018). Good Questions. In J. Dunn, & K. Ahlstrom-Vij (Eds.) *Epistemic Consequentialism*. Oxford, UK: Oxford University Press.

- Pettigrew, Richard (2020). Logical Ignorance and Logical Learning. *Synthese*.
- Pettigrew, Richard (2022). *Epistemic Risk and the Demands of Rationality*. Oxford, UK: Oxford University Press.
- Rinard, Susanna (2017). No Exception for Belief. *Philosophy and Phenomenological Research*, 94(1), 121–143.
- Rosenstock, Sarita, Bruner, Justin, & O'Connor, Cailin (2017). In Epistemic Networks, Is Less Really More? *Philosophy of Science*, 64(2), 234–252.
- Schoenfield, Miriam (2017). Conditionalization does not (in general) Maximize Expected Accuracy. *Mind*, 126(504), 1155–87.
- Schultheis, Ginger (ta). Accurate Updating. *Philosophy of Science*.
- Simion, Mona (2023). Resistance to Evidence and the Duty to Believe. *Philosophy and Phenomenological Research*.
- Staffel, Julia (2019). Credences and suspended judgments as transitional attitudes. *Philosophical Issues*, 29(1), 281–294.
- Staffel, Julia (2021a). Pro tem rationality. *Philosophical Perspectives*, 35(1), 383–403.
- Staffel, Julia (2021b). Transitional attitudes and the unmooring view of higher-order evidence. *Noûs*, 57(1), 238–260.
- Thorstad, David (2022). There are no epistemic norms of inquiry. *Synthese*, 200(410).
- Willard-Kyle, Christopher (forthcoming). The Knowledge Norm of Inquiry. *The Journal of Philosophy*.
- Williams, J. Robert G. (2018). Rational Illogicality. *Australasian Journal of Philosophy*, 96(1), 127–141.
- Williams, J. Robert G., & Pettigrew, Richard (2023). Consequences of Calibration. *The British Journal for the Philosophy of Science*.
- Williamson, Timothy (2013). Gettier Cases in Epistemic Logic. *Inquiry: An Interdisciplinary Journal of Philosophy*, 56(1), 1–14.
- Zollman, Kevin J. S. (2007). The communication structure of epistemic communities. *Philosophy of Science*, 74(5), 574–587.
- Zollman, Kevin J. S. (2010). The Epistemic Benefit of Transient Diversity. *Erkenntnis*, 72, 17–35.