Evaluating Inquiries

BPhil in Philosophy: Thesis



Trinity Term 2020

Word Count: 29,552

Pietro Cibinel

Abstract

How should we go about collecting evidence when we're looking to settle some question? In this thesis, I develop and defend an account of the evaluative and normative epistemic dimensions intrinsic to the practice of evidence-gathering. Chapter 1 introduces the topic. Chapter 2 is devoted to laying out the positive proposal, which makes use of the notions of accuracy and competence to spell out the conditions under which an act of evidence-gathering is successful, rational, and skillful. It also shows how the account can be applied to the simple case of inquiry into polar questions as well as to the more complex cases of inquiry into wh-questions and subject matters. Chapters 3 and 4 provide arguments for accepting the account. Chapter 3 zooms in on the central notion of the proposed account, *accuracy gain*, and argues that such a notion independently provides the best available measure of information value, namely the usefulness of a certain piece of information given the aim of settling a certain question. This is done by showing that accuracy gain fares better than competing measures from information theory, Bayesian statistics, and decision theory on a number of fronts. Chapter 4 first looks at how the proposed account fares *vis-à-vis* three main competing views about the aim of inquiry, arguing that it emerges as superior in several respects. Then, the Chapter considers three objections that may be raised against the proposed account, and argues

that these can be met. Chapter 5 concludes by looking at potential avenues for future research.

Acknowledgements

I would first like to thank my supervisor, Timothy Williamson, for his insightful comments on each of the chapters. The way I think about philosophy and in particular epistemology would be much different—crucially, poorer—if I hadn't had the pleasure of meeting him and studying his work. For the many philosophical conversations, which are one of the highlights of my BPhil, I thank Richard Roth and Sammy Boardman. I would also like to thank Aglaia von Götz, Daniel Kodsi, David Thorstad, Kevin Dorst, Matt Hewson, Richard Pettigrew, and Sebastian Liu for helpful written comments on parts of the thesis, and people at the Ockham Society and the Oxford Epistemology Reading Group for interesting discussions on the paper "Information Value and Accuracy Gain", which figures below as Chapter 3 of the thesis.

Contents

1 Introduction **5**

2 Evidence-Gathering: Accuracy and Competence **14**

3 Information Value and Accuracy Gain **42**

4 A Comprehensive Theory of Inquiry **72**

5 Conclusions **99**

Chapter 1

Introduction

This thesis begins with three simple stories:

> *The Detective.* At a summer mansion by the coast, there's been a murder. The detective has been called on the phone, and is already on her way to the crime scene. On the car, she thinks about how she should proceed. She needs evidence in order to find out who committed the crime. What she is thinking about right now is how best to collect such evidence. She knows that under interrogation, sometimes flaws can be found in the professed alibis of guilty individuals. But for that, one needs to ask the right questions. She also knows that finding the murder weapon would be a breakthrough in the investigation. But for that, one needs to look in the right places.
>
> *The Lab.* A group of scientists has just received a major grant, which allows them to pursue their research project. What they now need to do is gather as much data as possible. The data is collected via experiments, and the experiments need special equipment. Unfortunately, the grant can only cover the purchase of a limited amount

of such equipment. So, they sit down and discuss what experiments it would be best to pursue.

The Curious Layperson. José just watched a documentary which sparked in him an interest in the issue of whether human beings have free will. He decides to read a book to gain more information on the matter. He knows that this is a controversial topic on which people have held a variety of views, and that reading different books will plausibly lead him to formulate different thoughts on the topic. After looking at the front cover of a number of books, he is now thinking carefully about which one to buy.

These three cases share an important feature. The detective, the group of scientists, and the curious layperson all need to collect evidence so that they can get closer to, and possibly obtain, the truth about a question that they're interested in. They are inquiring into their respective matters, or are about to engage in such a process of inquiry. Further, they all reasonably ask themselves how best to pursue their inquiries. They know that their inquiries can go well or badly, better or worse, depending on what choices they make as to which sources of evidence to consult: they know that they can make good or bad, better or worse, choices of evidence-gathering. Cases of this kind will be the focus of this thesis. To abstract away from the specifics of the three cases above, I will be focusing on a rather idealized and simplified situation, which I call

The Setup. An agent is interested in settling a question, Q, which I will call the agent's *research question*. The agent can perform one of a set of mutually exclusive and jointly exhaustive acts of evidence-gathering: acts such that performing them results in the acquisition of certain pieces of evidence.

How should the agent choose? Under what conditions should we say that the agent chose the right act of evidence-gathering? Under what conditions should we say that the agent's inquiry went well? This thesis is concerned with the evaluative and normative dimensions intrinsic to the practice of evidence-gathering: with what makes certain inquiries go well and others badly, with what makes it the case that one ought to perform certain acts of evidence collection and not others.[1] Here's an initial answer to these questions, given for the purposes of circumscribing some more the boundaries of the topic of this thesis: an inquiry going well, and an agent choosing the right act, is to do with how informative, conclusive, telling, instructive, revealing, illuminating, the evidence that the agent collects is with respect to the agent's research question. This answer is non-committal and incomplete. First, it doesn't say what it is for some evidence to be informative, conclusive, telling, etc. Part of the thesis will be concerned with filling in these details. Secondly, it doesn't specify what the *exact* relationship is between the evaluative and normative judgements to be made about a given case and the evidence collected being informative, conclusive, telling, etc., in that case. Part of the thesis will again be concerned with filling in these details.

The evaluative and normative dimensions that I will be interested in are *intrinsic* to the practice of evidence-gathering in the sense that they hold independently of whatever other evaluative and normative considerations may be at play in each given case of inquiry. An agent may have, for example, *moral* or *prudential* reasons not to consult a certain source of evidence. In these ways, we can talk about inquiries going well or badly from a moral or prudential point

[1] As is standard in contemporary normative ethics, I think of the evaluative realm as having to do with value: with whatever makes things good or bad, better or worse. And I think of the normative realm as having to do with "oughts": with whatever makes acts right or wrong, permissible or impermissible. While I will end up endorsing a view which links the two very tightly at least with respect to evidence-gathering, it's important to keep them apart until any substantive points are made.

of view, and we can talk about which acts of evidence-gathering one ought to perform from a moral or prudential point of view. In this thesis, I will not be concerned with these kinds of evaluative and normative dimensions. I will be concerned with inquiries going well or badly *qua* inquiries, and with acts of evidence-gathering being the ones one ought or ought not to perform *given* the sole aim of settling one's research question. The detective, the group of scientists, and the curious layperson are not concerned with the moral or prudential upshots of their choices of evidence-gathering—at least as I've described the cases. When thinking about which questions to ask, which experiments to perform, or which books to read, their unique concern is settling their research question, or getting as close to its truth as their circumstances allow.[2]

For example, if the scientists choose to perform a set of experiments whose results disprove a number of false hypotheses about their question of interest, their inquiry will have gone well, and they will have chosen the right experiments to perform, granted that other experiments would have delivered more inconclusive data. If, by contrast, they choose to perform a set of experiments whose results fail to improve their understanding of the issue of interest, their inquiry will have gone badly, and they will have chosen the wrong experiment to perform, granted that other experiments would have delivered more conclusive data. These evaluations and normative judgements hold independently of any moral or prudential observations that might additionally be made about the case. I will have more to say about what makes such independent evaluative and normative domains possible for the practice of evidence-gathering; for now, my main point is to note that I will bracket moral, prudential, and other

[2]One may think: aren't the reasons why such agents are inquiring in the first place *prudential* reasons? Perhaps. But note that this doesn't threaten the existence of an evaluative and a normative dimension intrinsic to inquiry. It may well be, for example, that prudential reasons require that the agent consult a misleading source of evidence simply because the truth would be too harmful for the agent to bear. Nonetheless, this wouldn't count as an inquiry going well, and as the right choice of evidence-gathering to make, from the point of view of the evaluative and normative dimensions intrinsic to the practice of inquiry.

kinds of non-epistemic evaluative and normative considerations throughout this work.

Let me now zoom out for a moment and consider how the project that I will pursue in this thesis fits into epistemology as a whole. Any comprehensive epistemological theory should account for two central aspects of epistemic activity. One is *belief formation*. This aspect has been at the centre-stage of much of post-Gettier epistemological theorizing. The key questions to ask at this juncture are: *given* that a certain agent possesses such-and-such evidence, what doxastic attitudes should she form? What makes it the case that she should form precisely those doxastic attitudes in response to that evidence? What justifies belief? These are among the most fundamental questions in epistemology. But an overly narrow focus on them risks suggesting an incomplete picture of epistemic activity. For it shadows the fact that agents have a degree of control over what evidence they receive. We get much of our evidence via perception, simply by navigating our environment and without any conscious effort directed at collecting it. But sometimes we also decide to *look* for evidence bearing on a certain matter, in order to become *better* informed about it. Nor is this a feature peculiar to adult human beings: children and non-human animals also actively look for information, such as when a dog uses its olfactory capacities to individuate the location of a certain object of interest. Following Jane Friedman, we can call the former kind of evidence acquisition "accidental learning", and the latter kind "learning by design".[3] Thus, besides belief formation, epistemic activity involves a second central aspect, *evidence collection*. While belief formation concerns doing the best one can given the limited perspective on the world that one occupies, evidence collection concerns gaining a better perspective on the world in the first place.

[3]Friedman (forthcoming b, p. 6). Of course, as Friedman notes, some types of learning may fall in between these two extremes, which are to be seen as generating something like a continuum rather than a strict binary categorization.

My project sits squarely within this second aspect of epistemic activity, evidence collection, or learning by design. There is an important epistemological issue belonging to this second aspect that I will not be concerned with, however. This is the issue of whether there are *epistemic* obligations or duties to collect evidence in the first place. Rather, I look at the question of what choices an agent should make, and how well an agent's inquiry goes, *granted* that our agent has already decided to collect evidence on the matter. Whether the agent is under a duty to collect evidence—epistemic or otherwise—is orthogonal to the issue of whether her inquiry ends up going well, and of whether she makes the right choices given that inquiry's aim. I thus remain neutral about the issue of whether there are any duties to gather evidence, and I do not take my project to rely on their existence.[4]

As needs to be the case with a project of this length, I will make a number of epistemological assumptions in this thesis that I will not argue for. Fortunately, all of them have been defended compellingly by other epistemologists. It will be useful to list the central ones here, so as to make clear the theoretical framework within which the positive proposal of Chapters 2, 3, and 4, is to be understood. First, I assume that evidence consists of propositions.[5] This is standard, and needed to be able to make use of the tools afforded by probability theory and decision theory in epistemological theorizing. Secondly, I assume that if some proposition p is part of your evidence, then p is true.[6] No false propositions can be part of your evidence. This is not to be confused with the obvious truth that your evidence can make a false proposition likely. I will remain non-committal about what other conditions must obtain for a proposition to become part of

[4]See, e.g., Feldman (2000) and Cohen (2016) for scepticism about the existence of epistemic duties to gather evidence; and see, e.g., Worsnip (2019), Miracchi (2019), and Flores and Woodard (ms), for defences of the existence of epistemic duties or norms enjoining agents to gather evidence.

[5]See, e.g., Williamson (2000, pp. 194-200) for a defence of the view that evidence consists of propositions.

[6]A defence of this assumption can be found again in Williamson (2000, pp. 201-202), but also, e.g., in Littlejohn (2013).

your total body of evidence. For instance, a common view has it that p is part of your evidence *iff* you know p.[7] While all I say below will be compatible with this view about evidence, I will not commit to it. Thirdly, I assume that agents can lend various degrees of confidence to various propositions, and that these degrees of confidence are faithfully represented by the real numbers within the [0,1] interval, with 0 representing minimal degree of confidence, and 1 representing maximal degree of confidence. Namely, I assume that *credences* exist.[8] Further, I assume that ideally rational agents have credences whose respective numerical values conform to the probability axioms.[9] And I also assume that rational agents have credences whose strength conforms to the degree of support afforded by their evidence.[10] Next, I assume that the evaluative and normative dimensions intrinsic to the practice of evidence-gathering with which I'll be concerned belong to *epistemic* axiology and *epistemic* normativity. This means assuming a fairly broad conception of the epistemic realm. Such an inclusive conception is increasingly popular in recent epistemology.[11] Finally, on the methodological side, I assume that theories in epistemology are to be compared via the use of an abductive methodology which weighs up

[7] Williamson (2000, ch. 9).

[8] My proposal will thus obviously be compatible with the view according to which credences exhaust the range of doxastic attitudes that agents can take, for which see, e.g., Jeffrey (1970), Pettigrew (2016). But it will also be compatible with more moderate views according to which agents can additionally adopt the doxastic attitude of outright belief, for which see, e.g., Wedgwood (2012), Buchak (2014) (though, as will become clear in Chapter 4, some views about the relationship between outright belief and credence are better supported by my proposal than others). My proposal, however, will be incompatible with the kind of eliminativism about credences defended by, e.g., Holton (2014) and Horgan (2017).

[9] For pragmatic, Dutch Book-style, vindications of this claim, see Ramsey (1926) and de Finetti (1974). For non-pragmatic, accuracy-based, vindications of this claim, see especially Joyce (1998) and Pettigrew (2016).

[10] This position may be labelled "evidentialism", although it needn't share the internalist commitments found in Conee and Feldman (2004). We can use the term in Lasonen-Aarnio's sense, as the view that "rational or justified doxastic states are those that *track degrees of evidential support*"(p. 599, emphasis in original).

[11] See, e.g., Worsnip (2019) and Friedman (forthcoming a) for arguments in favour of such a broad conception of the epistemic realm. Those who adopt a more narrow conception of epistemic normativity and evaluations, e.g. Feldman (2000), may still find my proposal interesting, but they will think about the sort of normative and evaluative judgements that I will be concerned with as strictly speaking non-epistemic.

the degrees to which these theories possess theoretical virtues such as simplicity, explanatory power, consistency with our intuitions, connections with other disciplines, and so on.[12]

Before I move on to summarise the structure of the thesis, a final preliminary remark is in order. This is to do with the use I will make of some terms. So far, I've used "inquiry" and "evidence-gathering" or "evidence collection" interchangeably, and I will continue to do so throughout the thesis. I will also use "evidence" and "information" to mean the same thing: the true propositions that the agent acquires and then updates on in the process of inquiry (although, again, I remain non-committal about what this *acquiring* amounts to). I use such a variety of terms in an attempt to establish connections between different literatures in epistemology, philosophy of science, cognitive psychology, and statistics, all of which are interested in similar questions but share no common terminology. At any given point, I will tend to employ the terms that the literature with which I'll be interacting at that point makes use of. That scholars in these different literatures really have been using these different terms to mean the same things will become clear as we go along.

Let me now outline the plan for the rest of the thesis. As already noted, the positive proposal is developed and defended in Chapters 2, 3, and 4. Chapter 2 is devoted to laying out the account of the evaluative and normative dimensions intrinsic to evidence-gathering that I then defend in Chapters 3 and 4. In particular, Chapter 2 puts forward the conditions under which an act of evidence-gathering counts as successful, rational, and skillful. Crucial to this account will be the notion of *accuracy gain*, which corresponds to the change (positive, negative, or null) in the accuracy of the agent's evidential probability in the research question's true answer that updating on the evidence

[12]For a defence of the application of abductive methodology to philosophical theorizing, see, e.g., Williamson (2016).

collected by performing the chosen act brings about. Chapter 3 defends the account by showing that accuracy gain provides the best available model of *information value*, a notion that cognitive psychologists and statisticians have been interested in for the purposes of interpreting experiments about people's information-gathering practices. Chapter 4 defends the account further by linking it to the existing literature on the aim of inquiry, and by showing that it fares better than the theories that have been proposed in that literature. Chapter 5 concludes by looking at the potential for future research that the account unlocks.

Chapter 2

Evidence-Gathering: Accuracy and Competence

In this Chapter, I lay out my account of the evaluative and normative dimensions intrinsic to the practice of evidence-gathering. I will state the conditions under which an act of evidence-gathering counts as successful, rational, and skillful, and will further explore the connections between these three aspects. It takes little reflection to notice that evaluating people's evidence-gathering choices is not something confined to the epistemology classroom, but is rather one of the ways in which we ordinarily hold each other accountable *qua* intelligent beings. Here are some examples. We would think irrational a philosopher who tried to figure out whether body and mind are distinct by reading only Descartes' remarks on the matter. We would deem unsuccessful someone who wanted to gain information about whether there is water on Mars and found herself no wiser than before, after hours of careful research. And we would consider incompetent a detective who repeatedly failed to get any closer to the truth about the crimes they investigate. What follows in this Chapter is an

investigation of the structure and grounds of these kinds of evaluations.

2.1 Initial Remarks

It will be useful to start with some general observations about what kind of practice evidence-gathering is. First, evidence-gathering is an activity with an *aim*, or *goal*. This plausibly makes evidence-gathering an example of what Ernest Sosa calls critical domains of value: sets of interrelated entities, activities, or practices evaluated with respect to a central value, which is valuable for its own sake *relative to that domain*.[1] Consider as an illustration chess. A player's choices can be evaluated *vis-à-vis* the aim of chess, i.e. checkmating the opponent. We can say that the player did well or badly, that they should or shouldn't have moved a given piece. These evaluative and normative considerations are insulated from moral, prudential or aesthetic evaluations, and are made possible by the fact that chess is an activity with an aim, or goal. I will state the aim or goal of evidence-gathering neutrally as *improving one's epistemic position with respect to one's question of interest*, for example the question whether p. There is a conspicuous literature on what the aim or goal of inquiry is, which I will turn to consider closely in Chapter 4: knowledge, true belief and justified belief are the main contenders.[2] In that literature, the aim of inquiry is often stated in neutral terms as *settling* a certain question, or as *having* the answer to that question, for example the question whether p.[3] However, this formulation is too narrow for our purposes. There are cases in which there is no way for agents to settle whether p, but there are ways for them to gain more information on whether p; in such cases, agents may rationally carry out an act

[1] Sosa (2007, pp. 72-74). I am here following Kelp (ms) in suggesting that evidence-gathering or inquiry is a critical domain of value in Sosa's sense.

[2] Williamson (2000), Millar (2011), and Kelp (2014) argue for knowledge as the goal of inquiry; Kvanvig (2003) and Lynch (2004) argue for true belief; and Rorty (1995) and Davidson (2005) argue for justified belief instead.

[3] E.g., Kelp (2014), Kelp (ms).

of evidence-gathering in order to gain such information bearing on whether p.[4] Such acts can be evaluated *qua* acts of evidence-gathering just as those which aim at settling whether p can. Our neutral formulation of the aim of evidence-gathering covers both cases, where settling whether p counts as the special case of improving one's epistemic position with respect to p maximally.

Second, as emerges from this very last point, the aim of evidence-gathering is a *gradable* one. An act of evidence-gathering can not only fulfil its aim or fail to do so, but can succeed or fail with respect to its aim to different degrees. One may improve one's epistemic position with respect to p, but perhaps one could have improved it even more, had one collected a different set of evidence. Chapter 4 will discuss at length the advantages of thinking about the aim of evidence-gathering in gradable terms, as opposed to the more familiar binary terms discussed in the current literature.

Third, an act of evidence-gathering is an act in the thick sense of the term, in that it involves *choice* on the agent's part. In this respect, we should expect epistemic evaluations and norms of evidence-gathering to inherit some of the features of prudential and moral evaluations and norms, which equally concern choices. And on the other hand, we should expect epistemic evaluations and norms of evidence-gathering to be unlike epistemic evaluations and norms governing belief formation, a process which is plausibly not controlled by choice.

Given these three structural points about the practice of evidence-gathering, two methods appear particularly fit for investigating its evaluative and normative dimensions. That collecting evidence involves choice and a gradable aim suggests that we employ the tools of decision theory in our study. And that evidence-gathering is an activity with an aim in the first place, where success

[4]Consider, for instance, historical matters that just can't be conclusively settled due to the paucity of historical data.

and failure at least partly depend on the agent's dispositions, character traits, and skills, suggests that we couple decision-theoretic thinking with the tools of virtue theory. Instead of defending these methodological choices further, I put them to use in what follows and see what good they can do.

As the simple examples with which we started illustrate, it seems that we can evaluate acts of evidence-gathering across different dimensions, all of which concern, nonetheless, the epistemic domain. One may fail to improve one's epistemic position with respect to p by bad luck; alternatively, one could easily have figured out that p but didn't due to irrational choices; and so on. We should thus want an account of the axiology and normativity of evidence-gathering which does justice to its multi-faceted nature.

Putting the points so far together, I am ready to state the account that I turn to develop in detail in the rest of this Chapter. I propose that an act of evidence-gathering can be evaluated epistemically across three different dimensions, depending on whether it's successful, rational, and skillful. Two notions are deployed to account for these evaluative dimensions: *accuracy*—in particular, accuracy *gain*—and *competence*. The act is successful to the extent that it maximises accuracy gain; it is rational to the extent that it maximises expected accuracy gain; it is skillful to the extent that it manifests the agent's competence at collecting accuracy-maximising evidence. In turn, these three evaluative dimensions combine in different ways, the best case being one where the act of evidence-gathering is *skillfully successful*: it maximises accuracy gain and expected accuracy gain with respect to the question of interest to the agent thanks to the agent's competence at collecting accuracy-maximising evidence.

The plan for laying out and defending my *accuracy-based*, *virtue-theoretic* account goes as follows. I first present the account as it applies to the simple case of polar questions, also known as yes-no questions. In §2.2, I introduce

accuracy gain as the central notion at play in evaluations of acts of evidence-gathering. Then, I argue in §2.3 that the accuracy-based account needs to be supplemented with a virtue-theoretic story of how the agent's skill can affect the extent to which accuracy is achieved. In §2.4, I present the complete accuracy-based virtue-theoretic account for the case of polar questions. In §2.5 and §2.6, I extend the account to cover inquiry into wh-questions and subject matters: §2.5 sketches a theory of questions and subject matters, and §2.6 puts that theory to use in a generalized version of the accuracy-based virtue-theoretic account.

2.2 Accuracy

To make things simple, I will now focus on polar questions only, questions of the form whether p. Recall that our neutral gloss on the aim of evidence-gathering was: improving one's epistemic position with respect to p. An act of evidence-gathering thus *succeeds qua* act of evidence-gathering to the extent that it improves the agent's epistemic position with respect to p, and fails otherwise. In this sense, the act is evaluated positively to the extent that it succeeds, and negatively to the extent that it fails. What is the connection between the evaluative and the normative—between how well an agent's inquiry goes and what act the agent ought to perform? I shall propose that it is very simple, and that it can be explicated by conjecturing that the following epistemic norm holds:

Success-Neutral. If agent S gathers evidence as to whether p, S must choose act a such that a yields evidence on p which most improves S's epistemic position with respect to p relative to all acts available.

In this section, I defend a view as to what improving one's epistemic position with respect to p amounts to, thus taking a stance on what the aim of evidence-

gathering is, stating it non-neutrally. In doing so, I also lay out a formal framework for modelling choices between the acts of evidence-gathering that an agent may perform.

2.2.1 Success

To introduce my proposal, notice first that the results of collecting evidence about p impact directly on one's propositional justification for p, and only indirectly one's doxastic justification for p. Suppose that p, that my prior total evidence supports .5 credence in p, and that after I collect some evidence bearing on p, my posterior total evidence supports .7 credence in p. My epistemic position with respect to p has intuitively improved. But this is so only with respect to my *propositional* justification, for if I go on to form a .3 credence in p, discarding what my evidence supports, there's a sense in which my epistemic position with respect to p has worsened. This is the sense associated with *doxastic* justification. The case illustrates how what agents do with the evidence they receive is orthogonal to the evaluative and normative dimensions pertaining evidence-gathering, and rather concerns epistemic evaluations and norms on *belief formation*.

The foregoing underscores the point made in the Introduction that epistemic activity is composed of two central aspects: belief formation and evidence collection. We can now expand on that picture and present the following attractive connection between propositional and doxastic justification on the one hand, and collecting evidence and forming beliefs on the other: belief formation concerns making the most of the evidence (propositional justification) one has, by forming doxastically justified attitudes; evidence-gathering concerns getting better evidence (propositional justification) in the first place. Since evidence-gathering is to do with propositional justification, I will here be interested in the

agent's *evidential probabilities*, rather than the agent's actual credences, which may be irrational. The agent's evidential probabilities at a time are the probabilities conditional on the agent's total evidence at that time.[5] I take these to be of equal value as the credences propositionally justified by the agent's total evidence, i.e. the agent's evidential probability in p is x at time t *iff* the credence in p propositionally justified by the agent's total evidence is x at time t.

In light of this, we can advance our understanding of the aim of evidence-gathering: improving one's epistemic position with respect to p is to be cashed out in terms of getting better propositional justification or evidence for p. When is evidence good? Here's the answer I shall endorse: evidence is good to the extent that it pushes the agent closer to truth and bad to the extent that it pushes the agent closer to falsehood.

This thought can be stated more precisely by using the notion of *accuracy*. Roughly, the closer a probability function $P(\cdot)$ is to the ideal probability function—the one which assigns 1 to truths and 0 to falsehoods—, the more accurate $P(\cdot)$ is. If we think of $P(\cdot)$ as the credence function which is propositionally justified by one's evidence, or equivalently as one's evidential probability function, we have a more precise way to cash out the natural proposal. If p is true, evidence E is good with respect to p *iff* $P(p|E) > P(p)$. If p is false, evidence E is good with respect to p *iff* $P(p|E) < P(p)$. To formulate the proposal more precisely, some decision-theoretic machinery needs to be in place.[6]

Let $O = \{o_1, o_2, ..., o_n\}$ be the set of mutually exclusive and jointly exhaustive options available to the agent, representing the acts that the agent may decide to perform. And let $W = \{w_1, w_2, ..., w_n\}$ be the set of possible worlds. We suppose both these sets to be finite, at least for now. In standard decision

[5] I will not take a stance on the details of how the absolute prior probability distribution is determined; see Williamson (2000, ch. 10), Huemer (2009), Hedden (2015b), and Bradley (2020) for discussion.

[6] I use standard expected utility theory in what follows. However, much of what I say can be recast using a risk-weighted expected utility framework, as presented, e.g., in Buchak (2013).

theory, we then have a utility function U which takes an option o from O and a world w from W and returns a real number $U(o, w)$, measuring the utility of performing o at w. Finally, with $P(\cdot)$ defined as above, i.e. the agent's evidential probability function, the *expected* utility $ExpU_o$ of choosing option o is given by $\sum_{w \in W} U(o, w) \cdot P(\{w\})$.[7]

Following a recent but established tradition, we use decision theory to regiment the notion of accuracy.[8] Utility function U is replaced by a function A which measures accuracy, where $A(P(\cdot), w)$ measures the accuracy of $P(\cdot)$ at w. Since our focus is on cases in which agents are interested in whether p, we will sidestep issues arising when thinking about how to measure accuracy for credence *functions*, and think of A as measuring the accuracy of single credences instead—e.g., $A(P(p), w)$ gives the accuracy of $P(p)$ at w.

The standard way to measure accuracy is the so-called Brier Score.[9] Where v is proposition p's truth-value,

$$\textbf{Brier}(P(p), v) \coloneqq -(v - P(p))^2.$$

Some think that the Brier Score is *the* correct measure of accuracy: if that were the case, the accuracy A of $P(p)$ at w, $A(P(p), w)$, would equal $\textbf{Brier}(P(p), 1)$ in case p is true at w, and $\textbf{Brier}(P(p), 0)$ otherwise.[10] Others prefer to work with weaker assumptions, only laying out constraints which any plausible measure

[7] To simplify things, I assume act-state independence in what follows. When applied to the epistemic realm, which will be done presently, it's not clear how decision theory should deal with act-state dependence cases—see Caie (2013), Greaves (2013), and Konek and Levinstein (2019). I will here sidestep these issues, and work with an "unconditional" decision theory of the kind Savage (1954) proposed.

[8] For a representative sample of work in accuracy-first epistemology, see Joyce (1998), (2009), Greaves and Wallace (2006), Leitgeb and Pettigrew (2010), Easwaran (2013), Pettigrew (2013), (2016), Carr (2015), Bona and Staffel (2017), Staffel (2017), Schoenfield (2018), Horowitz (2019), Dorst (2019).

[9] The score was originally introduced by Brier (1950) as a way to assess forecasts.

[10] Pettigrew (2016, ch. 4) argues for the Brier Score as the uniquely legitimate measure of accuracy up to positive linear transformation. Additional arguments for the Brier Score are given, e.g., in Joyce (2009).

of accuracy should respect.[11] I will largely sidestep the issue of what measures of accuracy count as legitimate—my focus is on accuracy itself, however that is legitimately measured.

Now, back to evidence-gathering. We said that an act of evidence-gathering concerning p succeeds to the extent that it improves one's epistemic position with respect to p, which, we clarified, means providing one with better evidence on p. We also saw that new evidence is good to the extent that it improves accuracy, i.e. that it warrants credences which are more accurate than those warranted by the prior total evidence. Let us think of options in O specifically as *acts of evidence-gathering*.[12] Suppose o_1 yields evidence E_1 such that $A(P(p|E_1), w) > A(P(p), w)$. Then we can say that performing o_1 improves the agent's epistemic position with respect to p at w. More generally, let $\mathscr{A}$ be a function from O and W to $[-1, 1]$ such that for all k,

$$\mathscr{A}(o_k, w) = A(P(p|E_k), w) - A(P(p), w)$$

—where E_k is the evidence collected by performing o_k at w. We can think of $\mathscr{A}$ as measuring the "accuracy" of an act of evidence-gathering o with respect to p, intended as the difference between the accuracy score $P(p)$ gets upon conditionalization on whatever evidence the agent collects by performing o, and $P(p)$'s initial accuracy score. If $\mathscr{A}(o, w)$ is positive, performing o makes for an improvement in accuracy; if $\mathscr{A}(o, w)$ is negative, performing o makes for a decline in accuracy; and if $\mathscr{A}(o, w) = 0$, performing o makes no difference

[11] This is the standard working methodology; see, e.g., Schoenfield (2018), Dorst (2019), Campbell-Moore and Salow (forthcoming).

[12] Oddie (1997) and Campbell-Moore and Salow (forthcoming) also apply epistemic decision theory to evidence-gathering, for the purposes of determining whether it can ever be epistemically rational to avoid collecting evidence. My account to follow owes much to the groundbreaking connection that these papers establish between the accuracy-first literature and the literature on evidence-gathering stemming from Good (1966)'s theorem. Below, in Chapter 3, footnote 18, I discuss more thoroughly how my proposal relates to the frameworks that Oddie and Campbell-Moore and Salow make use of.

accuracy-wise. $\mathscr{A}$ measures what I will from now onwards call the *accuracy gain* that the evidence collected by performing the given act affords.

We are now in a position to replace **Success-Neutral** with the following:

Success. If agent S gathers evidence as to whether p at w, S must choose o from O such that for all k, $\mathscr{A}(o, w) \geq \mathscr{A}(o_k, w)$.

Informally, S must maximise accuracy gain, relative to all acts available.[13] We have now developed a picture of a first dimension of evaluation, and correspondingly a first normative dimension, governing evidence-gathering.

2.2.2 Rationality

Note, however, that one may conform to **Success** in a way which is irrational; and equally, one may rationally fail to conform to **Success**. To give an example of the first type of case, suppose a source which I have reason to think unreliable would in fact provide me with excellent evidence on p. It may be that consulting such a source is required by **Success**, if I am gathering evidence as to whether p. Still, it would not be rational for me to gather such evidence. To give an example of the second type of case, suppose that my evidence strongly suggests that looking into a given source would provide me with excellent evidence on p, but suppose also that my current evidence is misleading: if I were to consult such a source, it would tell me nothing as to whether p. It would be rational for me to consult such a source, but by doing so I would be violating **Success**.

These additional evaluative and normative dimensions are explained by reference to the distinction between the objective and subjective senses of "must",

[13] Notice that since I'm interested in the accuracy gain of the agent's *evidential probabilities* rather than that of the agent's actual credences, I'm not accepting a form of veritism about epistemic value. Or if I am, it is not the standard one: what is of final value with respect to evidence-gathering or inquiry is the accuracy of the agent's evidential probabilities rather than the accuracy of the agent's actual credences. Chapter 4 tackles this issue in detail.

"ought", and "should".[14] These deontic modal terms are highly context-sensitive.[15] In particular, uses of these terms are relative to salient bodies of information, or probability functions.[16] As "must" figures in **Success**, it is to be understood as relative to the omniscient probability function, or, equivalently, to the body of information which contains all and only truths. By contrast, in the two examples just given, what it would be rational for me to do is anchored to senses of "must", "ought", and "should" which are relativized to the body of information that I currently possess as evidence, or equally to the probability function justified by my current evidence.[17]

In the practical case, one is successful to the extent that one maximises utility, but one is rational to the extent that one maximises *expected* utility, the mathematical expectation of utility as calculated by means of the agent's evidential probabilities. To accommodate these additional evaluative and normative dimensions of evidence-gathering, we thus need to introduce the epistemic analogue of expected utility: expected accuracy. The expected accuracy of a certain credence in p, $C(p)$, relative to probability function $P(\cdot)$ and accuracy measure A is as follows:

[14]Philosophers who have argued for a distinction between objective and subjective "oughts" include: Ewing (1948), Brandt (1959), Parfit (1984), Jackson (1986), Jackson and Pargetter (1986), Gibbard (2005). Lasonen-Aarnio (forthcoming) has recently provided reasons for being sceptical about attributing any important role to subjective "oughts" in normative theory. I will not address Lasonen-Aarnio's arguments directly. However, it's important to note that if one were to employ Lasonen-Aarnio's framework in our context, one would not be able to make use of the tools of decision theory, which gives central importance to the notion of expected value and its corresponding subjective "ought" (or "should", etc.). Moreover, Chapter 3 in particular will show that the distinction between objective and subjective evaluations and norms can be put to much useful work in epistemology and cognate disciplines for the purposes of developing an account of information value. This will at least partly motivate its use in the account.

[15]See Åqvist (1984) and Lewis (1973) for expositions of the classical semantics for deontic operators.

[16]Cf. Björnsson and Finlay (2010) and Wedgwood (2017b). Björnsson and Finlay appeal to bodies of information, Wedgwood to probability functions.

[17]Of course, there will be even more subjective senses of "must", "ought", and "should", e.g. the one tied to the agent's *actual* credences, which may or may not conform to what the agent's evidence supports. However, for reasons mentioned above, I will here focus on the agent's evidential probabilities.

$$ExpA_{P(\cdot)}(C(p)) = P(p) \cdot A(C(p), p) + P(\neg p) \cdot A(C(p), \neg p).^{18}$$

Now, recall that we're interested in the values of $\mathscr{A}$, as these are the ones that determine which acts of evidence-gathering available to the agent are most successful. We can similarly devise the *expected accuracy gain* of an act of evidence-gathering o with respect to p, relative to $\mathscr{A}$ and probability function $P(\cdot)$ as follows:

$$Exp\mathscr{A}_{P(\cdot)}(o) = \sum_{w \in W} \mathscr{A}(o, w)P(\{w\}).$$

Suppose $Exp\mathscr{A}_{P(\cdot)}(o)$ is at least as high as the expected accuracy gain of any other act of evidence-gathering available to the agent. Intuitively, it is rationally permissible for the agent to choose o, even if o actually turns out not to maximise accuracy gain with respect to p. Or, suppose there are acts available to the agent which have higher expected accuracy gain than $Exp\mathscr{A}_{P(\cdot)}(o)$. Intuitively, it would be irrational for the agent to choose o, whether or not o actually maximises accuracy gain. This suggests the following:

Rationality. If agent S gathers evidence as to whether p, S's choice is maximally rational *iff* S chooses o from O such that for all k, $Exp\mathscr{A}_{P(\cdot)}(o) \geq Exp\mathscr{A}_{P(\cdot)}(o_k)$.

Informally, S should maximise expected accuracy gain—where "should" is to be interpreted subjectively, as expressing the requirements of rationality. One's (mathematical) expectations are widely assumed to determine what it is rational for one to do, even outside decision theory.[19] It seems that evidence-gathering

[18] As I do here, I will sometimes switch to using sets of worlds rather than single worlds as the second argument of the accuracy function, A, when no more fineness of grain is needed than that provided by the partition induced by the research question: here, "$A(C(p), p)$" is to be read as "$A(C(p), w_p \in p)$".

[19] See, e.g., Parfit (2011, p. 34), Hedden (2015a), Wedgwood (2017a).

is no exception in this respect. I use the expression "maximally rational" to underscore the point that the agent's act may display different degrees of (ir)rationality, just as it may display different degrees of success. For example, if the agent chooses the second-best act with respect to expected accuracy gain, her choice will be less than fully rational, but nonetheless not completely irrational. When talking about successful and rational acts *simpliciter*, I shall mean maximally successful and rational acts.

One may ask: surely what is rational to do does not depend on mathematical computations as complex as those laid out above, to which any human being has presumably at best limited access? The equations above, however, only unpack the mathematical structure of expected accuracy gain. This does not mean that we need to consciously compute anything when we choose how to act. Suppose I know that source of evidence α is very reliable, and that looking into it would provide me with extremely relevant information as to whether p. And suppose I know further that source β is quite unreliable, and that looking into it wouldn't tell me much of relevance as to whether p. Then, the expected accuracy gain of looking into α will be higher than that of looking into β, but I needn't compute anything to realize this. Chapter 3 will additionally suggest that expected accuracy gain can be maximised via use of certain readily available heuristics.

*

Summing up, we have identified two evaluative dimensions governing the practice of evidence-gathering, which are coupled by two corresponding epistemic norms. Do these suffice to account for the range of ways in which we can assess acts of evidence-gathering *qua* acts of evidence-gathering? In the next section, I argue that one final element is still missing.

2.3 Competence

Consider again the case just presented, where my evidence suggests that source α is reliable and relevant with respect to p, and that β is unreliable and not particularly relevant with respect to p. Suppose I go on to look into α, and collect the given evidence, thus maximising expected accuracy gain and, let us assume, accuracy gain too. I thus comply with **Success** and **Rationality** and my act is evaluated positively, as both successful and rational. However, suppose now that my policy was the following: toss a coin; if it lands heads, look into α, if it lands tails, look into β. Luckily, heads came up. Contrast this with the different case where I choose to look into α because I am responsive to the fact that looking into α is the best option avaliable given the aim of the practice I'm engaging in, namely evidence-gathering concerning p.

In both cases, I succeed and choose the rational option, but the first case should be negatively evaluated in a way that the second case should not. So a third dimension of evaluation seems needed to account for this evaluative difference between the two cases. What is required, I will suggest, is an appeal to the notion of *competence*, or *skill*. In the first case considered, while my act of evidence-gathering is successful and rational, it fails to be skillful, as it doesn't manifest any competence on my part at gathering good evidence, i.e. accuracy-maximising evidence. In the second case, by contrast, my successful and rational choice does seem to manifest a competence at gathering accuracy-maximising evidence.

How should we understand the notion of competence, or skill, especially as it relates to collecting evidence? Following a number of virtue epistemologists, I will understand competence or skill as a particular kind of *disposition*, namely a disposition to *succeed across a relevant range of scenarios*.[20] If we further think

[20]See especially Sosa (2015, ch. 4), but also Greco (2010, p. 77) and Beddor and Pavese (2018). By making use of some tools from virtue epistemology, I do not wish to commit to the claim that

of dispositions in modal terms, we get a modal account of competence, too: an agent S's performance is skillful to the extent that in a sufficiently high proportion of relevant worlds in which S engages in a sufficiently similar performance, S attains the aim of that performance.[21] While the modal account provides an informative way to think about competence, it is optional for the points I wish to make, and can be replaced with a different account by those who think that dispositions should not be given a modal treatment.

Given that the aim of evidence-gathering is maximising the accuracy of one's evidence, competence at evidence-gathering thus amounts to *a disposition to collect accuracy-maximising evidence across a relevant range of cases*. What determines which cases count as relevant? In general, relevance seems to be fixed by the interests of the community engaged in evaluating the given performance.[22] When assessing a goalkeeper's competence at stopping shots, for example, we're only interested in the kinds of scenarios which are likely to obtain in a game of football. Further, our interests may be such that we find it useful to draw additional distinctions: a keeper may be skillful at penalties but very bad in one-to-ones. The same remarks seem to apply to evidence-gathering. The range of cases relevant for competence at evidence-gathering will be determined by whatever categorisations best serve our interests as inquirers. For example, we may wish to group together cases in which agents inquire into a certain subject matter M, and single out a competence at collecting evidence about M. Then, when deciding whether the agent's act of evidence-gathering concerning p was skillful, where p belongs to subject matter M, we look at whether they were manifesting the specific competence at collecting evidence

knowledge is to be analysed via an appeal to epistemic virtues or competences. I simply borrow the helpful framework virtue epistemologists have put forward in thinking about performances in general, and leave it open whether knowledge itself is to be elucidated by means of this framework.

[21] Cf. Beddor and Pavese (2018).

[22] See again Sosa (2015, p. 104); cf. Lasonen-Aarnio (forthcoming), whose account of good dispositions is reliant on a similar notion of relevance.

about M, rather than a competence at collecting evidence more generally. We should thus make room for a variety of competences at evidence-gathering, determined by such things as subject matter, discipline, conditions of inquiry, and so on.

To sum up,

Competence. If agent S gathers evidence as to whether p, S's choice is skillful to the extent that it manifests the agent's relevant competence at gathering accuracy-maximising evidence.

Now we can ask: in what ways can competence (or its lack) influence my reaching (or failing to reach) the aim of evidence-gathering? I may lack competence and fail for this reason to collect accuracy-maximising evidence. Or, as in the first case of this section, I may display no competence but collect accuracy-maximising evidence nonetheless thanks to a stroke of good luck. Conversely, due to bad luck, I may manifest a competence at collecting good evidence and yet fail to be successful. If two strokes of luck are at play, I may be successful and competent, but not successful because competent. The case has this structure: I manifest my competence, but bad luck would make my act unsuccessful if it weren't for a stroke of good luck which makes it successful after all. Finally, my act can be successful *because of* my competence. In the next section, I will explore in more detail this special kind of case, after putting the pieces so far together in a presentation of the main account.

2.4 The Accuracy-Based Virtue-Theoretic Account

We can now state our complete account of the epistemic evaluations surrounding the practice evidence-gathering:

The Accuracy-Based Virtue-Theoretic Account.

- An act of evidence-gathering is *successful* to the extent that it maximises accuracy gain relative to all acts available;
- It is *rational* to the extent that it maximises expected accuracy gain relative to all acts available;
- It is *skillful* to the extent that it manifests the agent's relevant competence at gathering accuracy-maximising evidence;
- It is *skillfully successful iff* it maximises accuracy gain and expected accuracy gain thanks to the agent's competence at gathering accuracy-maximising evidence.

Two comments about the account are in order. First, while a skillful act of evidence-gathering may fail to be successful, it plausibly can't fail to be rational. In general, how could an agent manifest a competence at choosing the best option while failing to choose what *they themselves* should expect to be best in the given case? More specifically, at least with competences involving choice, it is hard to see how scenarios in which what the agent should expect to be best is in fact best could fail to be in the range relevant for competence-assessment: the range of counterfactual scenarios that we look at in order to decide whether the agent's act is skillful. If the competence is one at choosing the best option, it will always be relevant to see what happens at worlds in which expectations match reality. But notice that if the agent chooses irrationally, the disposition they manifest will lead them astray in such worlds.[23] Such a disposition will therefore not be one to succeed, and won't qualify as a competence. So there seem to be quite robust structural reasons for thinking that an irrational act of evidence-gathering can't be skillful: even if the irrational act succeeds at the actual world, the disposition that the act manifests will lead the agent to fail in

[23]This is for the same reason why we can't be disposed to wear seat belts only when car accidents occur—cf. Lasonen-Aarnio (forthcoming) on the notion of *feasible* dispositions.

those worlds in which what the agent should expect to be the best option is in fact the best option. If one thinks that some outlandish cases can be devised in which, though we're evaluating a competence at choosing the best option, we're not interested in worlds where expectations match reality, one can and should still accept the weaker claim that competence is good evidence for rationality.[24]

Secondly, skillfully successful acts of evidence-gathering occupy a special place in our account. Not only are they cases of evidence-gathering which get evaluated positively across all three dimensions, but they further constitute a species of the genus *success through competence*. These are performances that succeed at their aim *because of* the agent's competence.[25]

Now, what is it precisely for my success to be because of my competence? Taking a clue from virtue epistemology again, we might say that it should be the case that my competence figures as the *most salient* causal factor in the explanation of why my attempt is successful. If, with John Greco, we also hold that "[i]n the case of a deviant causal chain, salience goes to what is deviant, and away from what is normal or usual", we can conclude that salience in causal explanation is not only necessary but also sufficient for success to obtain because of competence.[26] A more modest proposal remains neutral about the relationship between salience and deviance and says that success is because of competence *iff* competence figures saliently and *appropriately* in the causal explanation of the success.

Alternatively, we might follow John Turri and Sosa and say that the agent's success should *manifest* the agent's competence.[27] If, with these authors, we further maintain that a success does not manifest a competence if the competence figures deviantly in the explanation of the success, we can say that success is

[24]That an act is rational, by contrast, isn't good evidence that it is skillful, as noted in §2.3.
[25]In Sosa (2007)'s terminology, these are *apt* performances.
[26]Greco (2010, p. 75)
[27]Turri (2011), Sosa (2015).

because of competence *iff* it manifests the relevant competence. A more modest proposal remains neutral about the relationship between manifestation and deviance and says that success is because of competence *iff* it manifests the relevant competence *appropriately*.

Whichever theoretical option we choose, I take it that we have a clear understanding of cases, and that we can reliably spot what counts as a competence appropriately accounting for the success. Bearing this in mind, I will stick to the above characterisation of skillfully successful acts, with the proviso that the *thanks to* relation is understood to exclude deviance. Evidence-gathering goes as well as it possibly can when agents rationally maximise accuracy gain thanks to their competence.

2.5 Partitions of Logical Space

Remember that so far we have focused on cases in which our agent is collecting evidence as to whether p. Such a narrow focus afforded a clearer outline of the main proposal. However, it is now time to consider the full range of scenarios in which agents engaged in a process of evidence-gathering may find themselves in. Beyond inquiring into polar questions, agents can inquire into wh-questions and subject matters. In this section, I sketch a theory of questions and subject matters, and introduce some important notions that this theory helps to elucidate. Then, in the final section of this Chapter, I put this theory to use to show how the account of §2.2-§2.4 can be extended to accommodate inquiry into wh-questions and subject matters.

Questions. Following an established tradition, I will model questions as *partitions of logical space*.[28] A question divides logical space into a set of mutually

[28]Hamblin (1958), Belnap and Steel (1976), Groenendijk and Stokhof (1984), Lewis (1988a). Versions of this theory are widely employed—see, e.g., Pavese (2017), Holguín (ms)—, though not without critics—see, e.g., Ciardelli et al. (2018).

exclusive ad jointly exhaustive cells, each corresponding to a complete answer to the question: the set of possible worlds at which that answer is true. Under this model, when trying to find the correct answer to a question, we're trying to figure out which cell the actual world is located in. Just as propositions have truth conditions, questions have *answerhood* conditions.

To give an example, the polar question *whether p* divides logical space in two cells: the one containing all and only p-worlds, the other contaning all and only $\neg p$-worlds. Wh-questions induce more fine-grained partitions. The question *who* ϕ divides logical space into as many cells as there are possible answers to the question: nobody ϕ, only John ϕ, only Mary ϕ, ..., only John and Mary ϕ, ..., everybody ϕ.

This way of thinking about questions allows us to shed light on a number of important notions. First, we can define partial answers to a question Q as unions of cells of the partition induced by Q. The more cells a partial answer excludes, the more informative it will be, if true. In the ideal case, an answer excludes all cells but the one containing the actual world: this will be the complete true answer to Q. Second, we can say what it is for a question to entail another, and for two questions to be equivalent. A question Q_1 entails another, Q_2, *iff* the cells of the partition induced by Q_2 are unions of the cells of the partition induced by Q_1. When a question entails another, settling the first will guarantee settling the second. If I learn who went to the party, I thereby know whether John did. Two questions are equivalent *iff* they induce the same partition. Third, we can represent the conjunction of two questions Q and R as the partition resulting from the pairwise intersection of the cells in the two partitions Q and R. The complete true answer to the question $Q \wedge R$ will be the intersection of Q's and R's complete true answers.

Subject Matters. I will follow David Lewis in thinking of subject matters

also as partitions of logical space.[29] For each cell in the partition induced by subject matter M, all worlds in that cell are exactly alike with respect to M, but may differ otherwise. For any two cells, each world in the first will differ from each world in the second with respect to M, and in exactly the same way. The actual world will be situated in the cell corresponding to how things actually are with respect to M. Thinking both of questions and subject matters as partitions of logical space blurs the difference between the two notions: each subject matter M induces the same partition as the question *what is the whole truth about M?*, and each question *who (where, what, ...)* ϕ induces the same partition as the subject matter *the truth about who (where, what, ...)* ϕ.

Because of this tight connection, we get notions for subject matters parallel to those set out above for questions. A proposition is entirely about a subject matter M *iff* it is the union of some cells of the partition induced by M. The more cells a proposition excludes, the more informative it will be if true. In the ideal case, a proposition excludes all cells but the one containing the actual world: this will be the whole truth about M. Second, we can say what it is for a subject matter to include another, and for two subject matters to be equivalent. A subject matter M_1 includes another, M_2, *iff* the cells of the partition induced by M_2 are unions of the cells of the partition induced by M_1. When a subject matter includes another, knowing the whole truth about the first will guarantee knowing the whole truth about the second. If I know all about reptiles, I thereby know all about crocodiles. Two subject matters are equivalent *iff* they induce the same partition. Third, we can represent the conjunction of two subject matters M and N as the partition resulting from the pairwise intersection of the cells in the two partitions M and N. The whole truth about $M \wedge N$ will be the intersection of M's and N's cells which contain the actual world.

[29] Lewis (1988a), (1988b). This theory is widely employed—see, e.g., Yalcin (2018)—, though not without critics—see, e.g., Yablo (2014).

In what follows, we will help ourselves of the partition model just outlined for the purposes of extending the *accuracy-based virtue theoretic* account to cover cases of evidence-gathering into wh-questions and subject matters. That the model turns out to afford a perspicuous and useful way to think about these issues will retrospectively motivate its employment.

2.6 The Simple Extension

The account developed in §2.2 states that collecting evidence concerning whether p is successful to the extent that the credence function propositionally justified by the new total evidence is more accurate *with respect to* p than the credence function propositionally justified by the prior total evidence. How can this thought be extended to cover cases in which one collects evidence in order to get closer to the truth of a wh-question and corresponding subject matter?

Here is the simple proposal I will put forward: inquiry into a question Q is successful to the extent that it maximises the accuracy of the complete true answer to Q. The account of inquiry into polar questions developed in §2.2-§2.4 then falls out as a special case of this more general proposal, which also accommodates inquiry into wh-questions and their corresponding subject matters. Consider the case where I'm a detective intending to find out who committed the crime. According to the account presently considered, my acts of evidence-gathering will be successful to the extent that the evidence I collect, together with my prior evidence, propositionally justifies a more accurate credence with respect to the complete true answer as to who committed the crime than the credence justified by my prior evidence alone.

Let me now delve into the details. Consider any wh-question or subject matter $Q = \{q_1, ..., q_n\}$, and let $q_w := \iota q \in Q : w \in q$, i.e. the complete true answer to Q at w. The success, or accuracy gain, of an act of evidence-gathering

o_k with respect to Q at w is measured as follows, where E_k is the evidence collected by performing o_k:

$$\mathscr{A}_Q(o_k, w) \coloneqq A(P(q_w|E_k), w) - A(P(q_w), w).$$

We are now in a position to state a more general version of the norm on successful evidence-gathering devised in §2.2.1, where the accuracy gain of any act of evidence collection is always relative to a question or subject matter Q:

Generalized Success. If agent S gathers evidence into Q at w, S must choose o from O such that for all k, $\mathscr{A}_Q(o, w) \geq \mathscr{A}_Q(o_k, w)$.

There are two ways for a piece of evidence E to make the complete true answer more likely than it was on the prior evidence. Evidence can exclude certain cells: E excludes cell q_i just in case $q_i \models \neg E$. The more cells in the partition induced by Q evidence E excludes at w, the more accurate $P(q_w|E)$ will be. When things go well at w, the evidence E collected excludes all cells but q_w. One's evidence will in such a case entail the correct true answer to the question Q, and $A(P(q_w|E), w)$ will be maximal. One will thereby be in a position to know the answer to the question. Alternatively, evidence can make some cells more likely than others, without ruling out any. In this way, E may make the complete true answer to Q very likely while assigning some very small probability to each other cell in the partition nonetheless. The account is neutral as to which of these two ways in which to improve accuracy with respect to q_w one should adopt, so long as the improvement in accuracy is the same. Only evidence which excludes cells can make one's accuracy maximal, however, and so whenever an agent can perform an act which would result in the collection of evidence which excludes all cells but the complete true answer to Q, the agent should perform that act.[30]

[30]This point holds at least with respect to the finite probability spaces with which I've been concerned so far; Chapter 4 will consider infinitary cases.

Given how §2.5 cashed out the notion of a conjunction of two questions or subject matters, our account easily accommodates cases in which an agent inquires into multiple questions or subject matters at the same time. If an agent looks for information to settle Q and R, we model her search as one for the complete true answer to the question or subject matter $Q \wedge R$. Her evidence-gathering will hence be successful just in case her posterior total evidence makes the cell in the partition induced by $Q \wedge R$ which contains the actual world as likely as possible.

Recall that in §2.2 we coupled the evaluative dimension of success with one concerning rationality, to account for the multi-faceted nature of the epistemic evaluations and normative judgements surrounding the practice of evidence-gathering. We can now provide a more general version of the principle stating the conditions for the rationality of an act of evidence-gathering, as this was set out in §2.2.2. Here is the natural extension: an act is rational to the extent that it maximises the expected accuracy gain of the agent's evidence with respect to the complete true answer to Q. More precisely,

Generalized Rationality. If agent S gathers evidence into Q, S's choice is maximally rational *iff* S chooses o from O such that for all k, $Exp_{P(\cdot)}\mathscr{A}_Q(o) \geq Exp_{P(\cdot)}\mathscr{A}_Q(o_k)$.

As far as the evaluative dimension pertaining to competence in concerned, no substantial amendment is needed when dealing with wh-questions and subject matters as opposed to polar questions. If the act manifests the agent's competence at gathering accuracy-maximising evidence, it will be skillful. And if it succeeds and is rational *thanks to* the agent's competence, then it will be skillfully successful. Extending our *accuracy-based virtue-theoretic* account to wh-questions and subject matters thus proves to be a relatively straightforward task.

2.7 Conclusion

This concludes the exposition of my account of the evaluative and normative dimensions intrinsic to the practice of evidence-gathering: success, rationality, and competence intertwine to explain the range of multi-faceted judgements that we can (and do) make about people's evidence-gathering choices. In the brief Interlude to follow, I explain what direction I will take in the next two core Chapters.

Interlude

Chapter 2 presented the *accuracy-based virtue-theoretic* account. The purpose of this brief interlude is to make the reader aware that the arguments of the next two Chapters will focus on defending primarily the *accuracy-based* component of the account. This is the component that stands most in need of support, and it is also the more distinctively epistemological one. That an evaluative dimension of competence is called for besides those of success and rationality emerges quite clearly from the examples of §2.3. What still needs defence is the claim that *accuracy gain*, as opposed to some other putative epistemic value, is required to account for the dimensions of success and rationality in the first place.

This is not to say that my virtue-theoretic interpretation of the examples of §2.3 cannot be questioned. For example, a more expansive theory of rationality in terms of reasons-responsiveness arguably accommodates much of the same data captured by virtue theory.[31] Consider again the case in which an agent chooses the rational option but only as a result of a coin toss. Such a reasons-responsiveness theory of rationality would re-describe the case as one where the agent chooses the rationally permissible act, but doesn't *act* rationally, because unresponsive to the reasons she has for choosing that option. This theory effectively includes (a kind of) competence as a necessary condition for rational

[31] See, e.g., Lord (2018) for a recent exposition of such a reasons-responsiveness theory.

action. Other normative theories distinguish between choosing the rational act, and being disposed to act rationally.[32] Such theories would re-describe the coin toss case as one where the agent chooses the rational act but fails to be disposed to act rationally. Since being disposed to act rationally is a kind of competence, these theories too arguably accommodate much of the same data captured by virtue theory.

The important upshot is that competence is at some level or other involved in the different ways in which virtue theory, reasons-responsiveness theories of rationality, and dispositional normative theories explain the examples of §2.3. Delving into detailed comparisons between these approaches would take us too far afield, away from epistemology and into normative theory and the theory of action. Given the space constraints for this thesis, I choose to focus on a thorough defence of the more distinctively epistemological aspects of the account developed in Chapter 2, namely those concerning the notion of accuracy gain. To be precise, let me then single out the component of the accuracy-based virtue-theoretic account that I set out to defend in Chapters 3 and 4 as follows:

> ***The Evidential Accuracy Theory of Inquiry:*** Inquiry is successful to the extent that the evidence that the agent collects *improves the accuracy* of the agent's *evidential probability* in the true answer to the agent's research question. Inquiry is carried out rationally to the extent that the evidence that the agent collects *is expected* to improve the accuracy of the agent's evidential probability in the true answer to the agent's research question.

The evidential accuracy theory of inquiry is nothing other than the conjunction of **Generalized Success**, **Generalized Rationality**, and the claim that the probability distribution $P(\cdot)$ relative to which accuracy is measured is the agent's

[32] See, e.g., Lasonen-Aarnio (2010), (forthcoming), and Williamson (2017) for expositions of such dispositional theories.

evidential probability distribution. The remainder of the thesis is devoted to arguing for the evidential accuracy theory of inquiry.

Chapter 3

Information Value and Accuracy Gain

In this Chapter, I provide an extended argument for thinking of *accuracy gain* as the central notion at play in epistemic evaluations and normative judgements surrounding the practice of evidence-gathering. To do so, I show that accuracy gain provides the best model of *information value*: the usefulness, informativeness, or value of a certain piece of information given the agent's aim of settling a certain question.

3.1 Measuring the Value of Information

Consider again what I called ***The Setup*** in the Introduction. We often find ourselves in situations just like that: we want to find out which of a set of mutually exclusive and jointly exhaustive hypotheses is true, and can perform a range of tests bearing on the matter. Some tests will be more helpful in settling the matter than others, since not all results will be equally informative. Here's an

extremely plausible principle about how we should act, independently of what has been proposed in Chapter 2: we should perform those experiments that maximise expected informativeness, and should be happy to the extent that the results we get are informative. But how are we to measure the informativeness of a given result in the first place? What is it for a piece of information to be valuable, or useful, given the aim of finding out the answer to a certain question? As in ***The Setup***, call the question whose true answer we are trying to find out our *research question*. And call the usefulness of information relative to a given research question the *value of information*, or *information value*. In this Chapter, I argue for the view that information value is to be understood in terms of *accuracy gain*.[1,2] Roughly, according to the proposed account, a piece of information is objectively informative or useful with respect to the research question to the extent that updating on it improves one's accuracy with respect to the true hypothesis; it is expectedly informative or useful with respect to the research question to the extent that updating on it is expected to improve one's accuracy with respect to the true hypothesis.

Most of the extant literature on the value of information is to be found within the field of cognitive psychology.[3] Cognitive psychologists, often drawing on previous work by statisticians, have proposed various measures of information value for the purposes of interpreting results about people's information-gathering behaviour, as reported by experimental studies on medical diagno-

[1]We should distinguish the notion of information value of interest here from Good (1966)'s notion, which concerns the extent to which acquiring information increases the maximum expected utility of a decision problem. Good's notion measures the *practical* usefulness of information, whereas the notion that I'm interested in measures the usefulness of information from a purely *epistemic* point of view. Oddie (1997) and Campbell-Moore and Salow (forthcoming) adapt Good's notion to the epistemic domain, but for purposes different from that of systematically laying out and defending an account of information value—cf. footnote 18.

[2]Cf. Roche and Shogenji (2018), who give an interpretation of the notion of mutual information from information theory in terms of inaccuracy reduction. I've developed my proposal independently of theirs—a thorough comparison between the two proposals will have to be postponed to future work. See footnotes 22 and 32 below for some further discussion.

[3]Nelson (2008) provides a helpful overview of this literature. Detailed references are provided below, as each account gets introduced.

sis, categorization, covariation assessment, and selection tasks.[4] Some of these measures have been employed as normative standards against which to evaluate the extent to which participants' choices in the above studies are rational; others as descriptive models of participants' behaviour. I will here argue that my new measure of the value of information, based on the notion of accuracy gain, is normatively superior to the measures that have so far been discussed by cognitive psychologists and statisticians. In §3.2, I present the accuracy gain model in detail, and then turn to argue for it in subsequent sections. Before that, however, it is necessary to acquire a clear view of the proposals already on offer, which I devote the rest of this section to.

Each of the models I shall presently outline is structured as follows. It first gives an equation for the usefulness, u, of an answer or result r relative to a research question $Q = \{q_1, ..., q_n\}$. Then, it defines the expected usefulness, eu, of a question or experiment R relative to the research question Q as the expectation of the usefulness of R's answers.

$$eu_Q(R) = \sum_j P(r_j) \cdot u_Q(r_j).^{5,6}$$

The four models of information value I will consider are probability gain, information gain, KL distance, and impact.[7]

[4]On medical diagnosis, see Baron et al. (1988); on categorization, see Skov and Sherman (1986) and Slowiaczek et al. (1992); on covariation assessment, see McKenzie (1994); on selection tasks, see Wason (1968), Klayman and Ha (1987), Oaksford and Chater (1994).

[5]Recall that throughout, $P(\cdot)$ is to be interpreted as the evidential probability distribution, or the credence distribution propositionally justified by the agent's total evidence, rather than the agent's (possibly irrational) *actual* credence distribution. When I talk about credences below, I thus intend these to be the credences propositionally justified by the agent's total evidence.

[6]In the equations to follow, I will leave it implicit that the usefulness of answers/results and questions/experiments is always relative to a research question, thus omitting the relevant subscript.

[7]I won't consider two other well-known measures, diagnosticity (Good (1950), Good (1975), Good (1983)) and log diagnosticity, because they are not suitable candidate measures for information value, as Nelson (2005) has persuasively argued.

3.1.1 Probability Gain

According to probability gain, a piece of information's usefulness is determined by the increase in the probability of making a correct guess (with respect to the research question Q) that updating on it effects.[8] We assume a utility function for the agent which gives the same positive utility to any correct guess and lower, equal utility to any incorrect guess. Under such conditions, the expected utility of making a certain guess is just its probability. The higher the probability of the most likely hypothesis, the higher the maximum expected utility. Thus, according to probability gain, the information value of answer r relative to Q is defined as follows:

$$u(r) = PG(r) = \max_{q_i} P(q_i|r) - \max_{q_i} P(q_i).$$

And the information value of question R relative to Q is:

$$eu(R) = EPG(R) = \sum_j P(r_j) \cdot PG(r_j).$$

3.1.2 Information Gain

According to information gain, a piece of information's usefulness is determined by the reduction in Shannon entropy (with respect to the research question Q) that updating on it effects.[9] Shannon entropy measures the uncertainty of a probability distribution in relation to a question's possible answers; it is maximal when each possible answer is assigned equal probability, and 0 when

[8] Baron (1985).

[9] Lindley (1956), Box and Hill (1967), Fedorov (1972), Oaksford and Chater (1994), van Rooy (2004).

all answers but one are excluded.[10] Here is how Shannon entropy is defined:

$$H(Q) = \sum_{q_i} P(q_i) \cdot \log \frac{1}{P(q_i)}.$$[11]

Suppose you acquire information r, and update on it. Then, the entropy of Q relative to your updated probability distribution may have changed, and can be recalculated as follows:

$$H(Q|r) = \sum_{q_i} P(q_i|r) \cdot \log \frac{1}{P(q_i|r)}.$$

The information gain model of information value equates the information value of answer r relative to Q with the reduction of entropy that getting r occasions:

$$u(r) = IG(r) = H(Q) - H(Q|r).$$

Then, the model defines the information value of question R relative to Q as the reduction of entropy that asking R is expected to occasion:

$$eu(R) = EIG(R) = \sum_{j} P(r_j) \cdot IG(r_j).$$

3.1.3 KL distance

According to Kullback-Liebler distance, a piece of information's usefulness is determined by the distance between the agent's prior probability distribution (with respect to the research question Q) and the probability distribution that updating on that piece of information generates.[12] The greater the belief change that getting the piece of information causes, the more useful such a piece of

[10] Shannon (1948).

[11] The convention that $0 \cdot \log 0 = 0$ is standard in information theory and easily justified by continuity arguments—see, e.g., Cover and Thomas (2006).

[12] Kullback and Leibler (1951), Cover and Thomas (2006).

information will be considered according to this measure. Formally, on the KL distance measure, the information value of answer r relative to Q is:

$$u(r) = KL(r) = \sum_{q_i} P(q_i|r) \cdot \log \frac{P(q_i|r)}{P(q_i)}.$$

And the information value of question R relative to Q is:

$$eu(R) = EKL(R) = \sum_{j} P(r_j) \cdot KL(r_j)$$

3.1.4 Impact

According to impact, a piece of information's usefulness is determined by the absolute change from prior to posterior probability distribution (with respect to the research question Q) that updating on the piece of information brings about.[13] Thus, impact agrees with KL distance that usefulness of information is to be cashed out in terms of belief change, but measures belief change in a different way. In particular, the impact model defines the information value of answer r relative to Q as:

$$u(r) = I(r) = 0.5 \cdot \sum_{q_i} |P(q_i|r) - P(q_i)|.$$

And it defines the information value of question R relative to Q as:

$$eu(R) = EI(R) = \sum_{j} P(r_j) \cdot I(r_j).$$

*

[13]Wells and Lindsay (1980), Klayman and Ha (1987), Nickerson (1996). These authors define impact in relation to polar research questions only; the more general formulation here employed is due to Nelson (2008).

Probability gain, information gain, KL distance, and impact largely agree about which answers and questions count as most useful in particular cases. However, circumstances can be devised in which they provide competing verdicts.[14] We can compare the models along two different dimensions; we can ask:

- Which of these models provides the best *descriptive* account of people's information-gathering practices?
- Which of these models provides the best *normative* account of how people ought to conduct their information-gathering practices?

Cognitive psychologists disagree about the extent to which an answer to the first question constrains an answer to the second. Some think that if a model doesn't meet basic predictive standards, it shouldn't be taken seriously as a normative proposal either; others allow for more discrepancy between the two levels.[15] It is important to note that the question with which I started the paper, and which will mainly concern the rest of it, is the *normative* one. When we ask what information value truly amounts to, we're asking how people ought to collect information, epistemically speaking—at least granting the principle stated at the start that rational people maximise expected information value, and successful people maximise information value itself.

In the following section, I lay out the details of a novel measure of information value, which I call the *accuracy gain* measure. Then, in subsequent sections, I argue for the superiority of accuracy gain as a normative standard over the extant measures just outlined.

[14] Nelson (2005).

[15] See, e.g., Oaksford and Chater (2007) for an approach sympathetic to the former view and, e.g., Baron (1985, pp. 52-53) for remarks gesturing towards the latter.

3.2 Accuracy Gain

I will use the Brier Score as a measure of accuracy in what follows, but this is mostly for illustrative purposes.[16] As already noted, my focus is on accuracy itself, however that is legitimately measured.

Suppose I'm interested in settling a research question Q, and collect some information r. Updating on r may alter my evidential probability distribution over Q's possible answers. Letting q_T be Q's true answer, we can ask whether updating on r resulted in an improvement in the accuracy of my evidential probability in q_T. If that is so, my epistemic position with respect to my research question Q has intuitively improved; if that isn't so, then it intuitively hasn't. Regimenting this thought yields **Objective Information Value = Accuracy Gain**:

(OIV = AG). The objective information value of r is $AG(r) = A(P(q_T|r), q_T) - A(P(q_T), q_T)$.[17]

Whether my accuracy has improved is something that can only be told by someone who knows what the true answer to Q is, a position which I would only occupy in the most favourable case where r excludes all possible answers but q_T. Thus, OIV = AG should be coupled with an account of *expected* information value, which is naturally cashed out as the accuracy gain that updating on r is *expected* to bring about. We thus get **Expected Information Value = Expected Accuracy Gain**:

(EIV = EAG). The expected information value of r is $EAG(r) = \sum_{q_i} P(q_i|r) \cdot$

[16]In particular, my main arguments, in §3.3, §3.4, and §3.5, do not rely on this choice: that is, they go through with any plausible measure of accuracy satisfying standard constraints, especially strict propriety. The choice does affect the points I develop in §3.6, but similar results can likely be achieved using alternative measures of accuracy.

[17]As I did at one point in Chapter 2, I am here using sets of worlds rather than single worlds as the second argument of the accuracy function, A, because no more fineness of grain is needed than that provided by the partition induced by the research question: "$A(P(q), q)$" is to be read as "$A(P(q), w_q \in q)$".

$$[A(P(q_i|r), q_i) - A(P(q_i), q_i)].$$

The models presented in §3.1 are concerned with a notion of information value *simpliciter*, and do not mark the difference that the present model draws between objective and expected information value. That the accuracy gain model draws this difference will play a crucial role in the arguments below; for now, however, let me simply note that it is *expected* accuracy gain which lines up with the notion of information value that the models of §3.1 attempt to account for, rather than objective accuracy gain. Even when accounting for the information value of single answers, the models of §3.1 are concerned with usefulness as this is perceived from the agent's point of view, and not by an external observer acquainted with the true answer to the agent's research question Q.

To see this, consider the case where an agent gets a piece of information r that makes one false answer very likely on their posterior evidence, and hence every other answer—including the true one—unlikely. All of probability gain, information gain, KL distance, and impact agree that r is very useful relative to Q (at least assuming that the agent's prior probabilities were broadly distributed), and the accuracy gain model agrees too, insofar as information value *simpliciter* is cashed out in terms of expected information value. Of course, though, since the actual accuracy of the agent's evidential probability in the true answer decreases, r has negative objective information value according to the present account. Since I am interested in comparing the accuracy gain account with the four extant accounts presented in §3.1, it will come helpful to place all accounts on the same level, and treat EIV = EAG as the "core" of the accuracy gain account, specifying the measure for information value *simpliciter*. As noted, however, the notion of objective information value will be central to the arguments to follow.

We can next define the expected information value of a question, R, relative

to a research question Q in the standard way, as the expectation of the (expected) information value of R's answers:

$$EIV(R) = EAG(R) = \sum_j P(r_j) \cdot EAG(r_j).^{18}$$

While it won't be as useful as in the case of answers, for completeness we can also define the objective information value of question R as the objective information value of its true answer. This concludes my exposition of the accuracy gain model of information value. Although it hasn't received any attention in the cognitive psychology literature on information value, I will now argue that there is much to recommend it.

3.3 Nonnegativity and Additivity

In this section, I show how the accuracy gain account solves a problem concerning synchronic vs. sequential learning. One illuminating way of getting a purchase on the difference between the various measures of information value here presented is to sort them out according to whether they are nonnegative or additive. A *nonnegative* measure always assigns positive information value to information which changes one's credences. An *additive* measure is such that for all pieces of information r_1, r_2, $u(r_1) + u(r_2|r_1) = u(r_1 \wedge r_2)$. On an additive measure, two answers' sum utility is the same whether the answers are given

[18]Campbell-Moore and Salow (forthcoming), in the course of showing that if Buchak (2013)'s risk-avoidance is rational, then it can be epistemically rational to refrain from collecting new evidence, define the (risk-weighted) expected accuracy of an act of evidence-gathering as the (risk-weighted) expectation of the accuracy of the credence you'd end up with after updating on the evidence collected. Oddie (1997) also measures the expected cognitive value (i.e. expected accuracy) of an act of evidence-gathering similarly. However, accuracy *gain* seems a better measure in the context of providing an account of information value. Suppose you have a prior justified credence of .99 in an answer to Q, and can ask a question R whose answers you know incapable of affecting your credence about the answer to Q substantially. If it's expected accuracy *simpliciter* that measures information value, asking question R has great expected information value, since its expected accuracy is high. Only using expected accuracy *gain* as a measure of expected information value gives the welcome result that the expected information value of R is low.

all at once or sequentially.

Probability gain and information gain are additive but can be negative, whereas KL distance and impact are nonnegative but not additive. Here is the problem: no measure can satisfy both nonnegativity and additivity, and yet failures of each principle lead to what seem like implausible results.[19] With respect to failures of nonnegativity: how can getting new information ever be of negative value in relation to one's aim of settling a given research question? With respect to failures of additivity: how can the order in which you receive two pieces of information matter with respect to their usefulness?

I will now discuss each problem in turn, illustrating with examples. To reveal my cards: I will agree with, and build on, extant criticism of additive measures, and argue that the standard criticism of nonnegative measures is misplaced. However, I will then show that each of nonnegativity and additivity has a role to play in a satisfactory account of information value, and that the accuracy gain model is uniquely placed to accommodate for this.

3.3.1 Nonnegativity Failures

Let me start with additive measures: probability gain and information gain. Suppose the polar question whether p is your research question. Your prior evidential probabilities are as follows: $P(p) = .3$, $P(\neg p) = .7$. You then get information r, which has the following evidential import: $P(p|r) = P(\neg p|r) = .5$ According to both probability gain and information gain, getting r has negative usefulness, or negative information value, with respect to your research question. Evans and Over (1996) rightly point out that this is an odd conclusion.

[19]Proof: suppose our agent is interested in whether p, with prior $P(p) = x$. Consider the pieces of information r_1 such that $P(p|r_1) = x' \neq x$, and r_2 such that $P_{r_1}(p|r_2) = x$. Now, consider a nonnegative measure of information value NIV. By nonnegativity and the fact that $x' \neq x$, $NIV(r_1) > 0$ and $NIV(r_2|r_1) > 0$. So $NIV(r_1) + NIV(r_2|r_1) > 0$. But $NIV(r_1 \wedge r_2) = 0$, for $P_{r_1}(p|r_2) = x = P(p)$. So, for any nonnegative measure NIV, in this class of cases $NIV(r_1) + NIV(r_2|r_1) \neq NIV(r_1 \wedge r_2)$, i.e. nonnegativity entails violations of additivity.

It is worth citing their contribution in full; while they argue against Oaksford and Chater (1994)'s information gain account, their point can be extended to any measure which can be negative:

> [...] if [$P(p)$] actually goes up from .3 to .5, after we gather some evidence, then we apparently lose information by gaining uncertainty. Of course, someone could just have the goal of reducing uncertainty in the Oaksford and Chater sense; but most people who have a hypothesis will surely have some interest in finding its truth value, and they should not use Oaksford and Chater's measure as a normative standard.[20]

They say next that in the case described "[...] we might suffer a cost in disappointment, but we would think of ourselves as having had an epistemic benefit".[21]

To expand on Evans and Over's criticism, we should look for a principled epistemological story backing up their intuition that people interested in finding the truth about their research question shouldn't use additive measures, which can be negative.[22] To this end, recall first that probability gain and information gain, just as any other measure presented in §3.1, are concerned with usefulness of information as this is perceived from the agent's point of view, and not by an external observer already acquainted with the true answer to the agent's research question. That the value of r is *negative* in our case, in virtue of the fact that it shifts $P(p)$ from .3 to .5, must then mean that the agent is evaluating the .5 credence from the perspective of the .3 credence. The message encoded in such a negative assignment of value is: *I am moving further away from the truth*. And it is precisely the other way around that any plausible measure of

[20] Evans and Over (1996, p. 358).

[21] Ibid.

[22] Roche and Shogenji (2018) also rehearse Evans and Over's criticism, but they too limit their arguments to an appeal to intuitions, for which they lack a principled epistemological foundation.

information value should evaluate such credences.[23] *Before* the agent gets r, of course, the .5 credence will be seen as sub-optimal from the agent's perspective, which is that of the .3 credence. But *after* r is acquired, the agent's point of view becomes that of the .5 credence, and thus it's the .3 credence which should be evaluated from the perspective of the .5 credence, contrary to what additive measures predict. A correct assignment of information value to r should encode the message: *good luck I got r, I guess my prior evidence supported being* too *confident in* $\neg p$.

At any given time, it is the credence justified by the agent's evidence which constitutes the point of view from which the agent should evaluate the usefulness of information, and of adopting other credences. If once one learns r, r becomes part of one's evidence, we have an explanation for why getting r can never be rationally considered of negative information value or usefulness when information value is assigned from the agent's point of view. Measures which can be negative, such as probability gain and information gain, violate the principle that the credences justified by the agent's evidence are always *self-recommending* in the way just laid out.[24] Since they entail such an implausible epistemological theory, additive measures should be rejected.

3.3.2 Additivity Failures

What about the costs of accepting a nonnegative measure? Consider the following case. Your prior justified credences about the possible answers to your polar research question are again $P(p) = .3$ and $P(\neg p) = .7$. Further, consider the following two pieces of information, r and $r*$, such that $P(p|r) = .7$ and

[23] As Evans and Over note, proponents of information gain may say that they're modeling agents purely interested in reducing uncertainty, in which case there is nothing wrong with the encoded message: *I am moving away from certainty*. But then this wouldn't be a model applicable to *inquiring* agents, for inquiry is directed at truth, not certainty. At any rate, proponents of information gain use this measure precisely in contexts in which agents aim at finding out the truth about a given question—see, e.g., Lindley (1956, p. 987), Oaksford and Chater (2007, p. 170).

[24] There are accuracy-theoretic tools for making this principle more precise: see footnote 26.

$P(p|r \wedge r*) = .3$. According to measures which violate additivity, such as KL distance and impact, if you learn r and $r*$ simultaneously, the information value of $r \wedge r*$ is 0, since gaining such information does not change your prior justified credences. If, however, you first learn r, and subsequently learn $r*$, KL distance and impact regard each piece of information as useful, despite the fact that your credences at the end of the learning process are the same as those with which you started out. In both cases, you start with the same prior credences, end with the same posterior credences, yet value the information differently. Some take this to be a counterintuitive result, a "strange contradiction".[25]

The result can certainly be made to *sound* counterintuitive—but how bad is the epistemological theory which licenses it? Once we examine the case carefully, it appears that there is nothing epistemologically flawed about such a differential assignment of information value. The action lies in the sequential case, so let's consider each step in the succession closely. We agreed that all measures of information value presented in §3.1, including KL distance and impact, assign usefulness as perceived from the agent's point of view, which we noted above corresponds to the agent's total body of evidence. So suppose your total evidence at the start of the sequential learning process is E. You first learn r, and your evidence becomes $E \wedge r$. The value of information r is a function of your prior point of view, E, and your current, and thus privileged, point of view $E \wedge r$. You evaluate r as useful. You then learn $r*$, and your evidence becomes $E \wedge r \wedge r*$. The value of information $r*$ is a function of your prior point of view, $E \wedge r$, and your current, and thus privileged, point of view $E \wedge r \wedge r*$. You evaluate $r*$ as useful. So far so good.

The alleged problem lies in the fact that you would have evaluated $r \wedge r*$ as useless had you learned both pieces of information simultaneously. Consider: you learn $r \wedge r*$, and your evidence becomes $E \wedge r \wedge r*$. The value of information

[25]Nelson (2008, p. 153).

$r \wedge r*$ is a function of your prior point of view, E, and your current, and thus privileged, point of view $E \wedge r \wedge r*$. So:

- When you evaluate r as useful you do so from the perspective of evidence $E \wedge r$, coming from the perspective of E;
- When you evaluate $r*$ as useful you do so from the perspective of evidence $E \wedge r \wedge r*$, coming from the perspective of $E \wedge r$;
- When you evaluate $r \wedge r*$ as useless, you do so from the perspective of evidence $E \wedge r \wedge r*$, coming from the perspective of E.

In no two such cases do you evaluate the same piece, or sets of pieces, of information as both useful and useless coming from the same old perspective and occupying the same new perspective. That would be an unwelcome result, a "strange contradiction", but it's not one delivered by nonnegative measures such as KL distance and impact. Since the pieces of information r and $r*$ are evaluated relative to different combinations of old and new perspectives, it is only to be expected that they get assigned different information values in the synchronic and sequential case. Shifts in perspective explain additivity failures. The theory of synchronic and sequential learning underwritten by nonnegative models of information value presents no particular challenge.

3.3.3 Accuracy Gain, Nonnegativity, and Additivity

Yet, can there really be no cases in which getting new information hinders one in one's attempt to get closer to the truth about one's research question? For example, an agent who gets some misleading evidence *does* seem to be worse off relative to her aim of settling the research question. And shouldn't the usefulness of information, at least in some sense, be independent of the order in which it is received? Although the alleged problem with nonnegative

measures can be defused, and, by contrast, additive measure are shown to commit to a flawed epistemology, the additivity principle does seem to latch onto something important about information value. Sticking with KL distance or impact doesn't get us any further in trying to understand what good there is about additivity.

How does the accuracy gain model fare in this debate? The measure of information value defined by EIV = EAG is nonnegative.[26] Thus, it avoids the problematic epistemological commitments of additive measures. However, it gets us no further in our attempt to understand the intuitive appeal of additivity, since it involves violations of this principle, just as KL distance and impact do.[27] But recall now that EIV = EAG is only part of the accuracy gain model of information value, accounting for what I called *expected* information value, that is, information value as evaluated from the agent's perspective—which is how information value is standardly understood in the cognitive psychology literature. The accuracy gain model, however, also offers a notion of *objective* information value, through the principle OIV = AG. And *this* measure turns out to be additive.[28]

[26]This is because any plausible measure of accuracy, such as the Brier Score, will be strictly proper: the expected accuracy of a credence x in p relative to a probability distribution on which $P(p) = x$ will be higher than that of any other credence—i.e. each probability distribution will take itself as giving the "best shot at the truth". Once the agent updates on new information r, thus, the expected accuracy of the posterior probability distribution, evaluated from the perspective of the posterior probability distribution itself, will be higher than that of the prior probability distribution. This means that the expected accuracy gain will be positive, unless there is no change in credence, in which case it will be 0.

[27]As an illustration, consider the case Nelson (2008) uses in relation to KL distance and impact. Your prior justified credences about the polar research question's possible answers are $P(p) = .75$ and $P(\neg p) = .25$. Now, consider information r_1 such that learning r_1 brings $P(p)$ down to .25, and information r_2 which, in conjunction with r_1, brings $P(p)$ back to .75. Since the credences don't change if you get both r_1 and r_2 simultaneously, $EAG(r_1 \wedge r_2) = 0$. However, $EAG(r_1) = .25 \cdot (-.5) + .75 \cdot (.5) = .25$. And $EAG(r_2|r_1) = .75 \cdot (.5) + .25 \cdot (-.5) = .25$. So, $EAG(r_1) + EAG(r_2|r_1) = .5 \neq EAG(r_1 \wedge r_2)$.

[28]Proof that for all r_1, r_2, $AG(r_1) + AG(r_2|r_1) = AG(r_1 \wedge r_2)$: by the definition of accuracy gain, $AG(r_1) + AG(r_2|r_1) = [A(P(q_T|r_1), q_T) - A(P(q_T), q_T)] + [A(P_{r_1}(q_T|r_2), q_T) - A(P_{r_1}(q_T), q_T)] = A(P_{r_1}(q_T|r_2), q_T) - A(P(q_T), q_T)$. By the same token, $AG(r_1 \wedge r_2) = A(P(q_T|r_1 \wedge r_2), q_T) - A(P(q_T), q_T)$. So, if $P_{r_1}(q_T|r_2) = P(q_T|r_1 \wedge r_2)$, then $AG(r_1) + AG(r_2|r_1) = AG(r_1 \wedge r_2)$. And of course, by the ratio formula: $P_{r_1}(q_T|r_2) = P_{r_1}(q_T \wedge r_2)/P_{r_1}(r_2) = P(q_T \wedge r_2|r_1)/P(r_2|r_1) = P(q_T \wedge r_1 \wedge r_2)/P(r_1)/P(r_1 \wedge r_2)/P(r_1) = P(q_T \wedge r_1 \wedge r_2)/P(r_1 \wedge r_2) = P(q_T|r_1 \wedge r_2)$. So, by modus ponens: $AG(r_1) + AG(r_2|r_1) = AG(r_1 \wedge r_2)$. We need not worry about conditionalization on propositions

While expected information value is a measure of how useful a piece of information is to the agent according to the agent's own perspective, objective information value is a measure of how useful a piece of information is to the agent according to an external observer already acquainted with the truth about the agent's research question. Since a piece of information may increase but also decrease the accuracy of the agent's evidential probability in the true answer to the research question, accuracy gain as a measure of objective information value can be negative. But this is exactly what we'd want out of an *objective* measure of information value. As noted earlier, if you get information which is misleading relative to your research question, you're moving away from the truth: such information has, objectively speaking, negative usefulness for you. While it is implausible that some information could ever be of negative value as judged from your own perspective—which was the problem with additive measures of, we can now say, *expected* information value—information is sometimes of negative value when judged from an external perspective.

Further, we can also see that there *is* a sense in which the usefulness of two pieces of information is independent of whether they are acquired simultaneously or sequentially; it is the sense associated with objective information value. And we can clearly see the intuition behind this thought: surely, an external observer acquainted with all the facts wouldn't be sensitive, in her judgement about how useful certain information is to you, to whether such information is acquired all at once or sequentially. Recall that additivity failures are generated by shifts in the agent's perspective. But the perspective of an external observer remains fixed throughout your information-acquisition process: no shift, no failure of additivity. The accuracy gain model of information value is

which have 0 prior probability, for (i) the propositions which the agent conditionalizes on constitute evidence and thus must be true, and (ii) $P(\cdot)$ represents the credence distribution justified by the agent's evidence, which never gives probability 1 to a false proposition (at least in the finite cases with which I've been concerned so far).

thus uniquely suited to account for both nonnegativity, which is deal with at the expectational level, and additivity, which is dealt with at the objective level.

3.4 Surprise, Excluded Possibilities, and Belief Change

In this section, I strengthen the case for the accuracy gain model by showing that it avoids additional problems faced by its contenders. In particular, I show that accuracy gain is the only model capable of accounting at once for the value of surprise, of excluding possibilities, and of changing beliefs.

3.4.1 Surprise

In this subsection, I show that information gain fails to account for the value of surprise. Suppose your research question Q has four possible answers, q_1, q_2, q_3, q_4. Suppose further that these are your priors: $P(q_1) = P(q_2) = .4$, $P(q_3) = P(q_4) = .1$. Consider next a counterpart of yours who is interested in a research question, $Q*$, with the same number of possible answers, and whose priors are $P'(q*_1) = P'(q*_2) = .4$ and $P'(q*_3) = P'(q*_4) = .1$. Suppose you get information r which excludes q_3 and q_4, so that your posterior probabilities are $P_{new}(q_1) = P_{new}(q_2) = .5$. Your counterpart, by contrast, gets information $r*$ which excludes $q*_1$ and $q*_2$, so that their posterior probabilities are $P'_{new}(q*_3) = P'_{new}(q*_4) = .5$. Intuitively, $r*$ is more useful to your counterpart than r is to you. Information r only confirms what you already thought was likely, namely that the true answer to Q is either q_1 or q_2. By contrast, information $r*$ not only excludes two possible answers, as r does, but remedies the negative effects of some prior misleading evidence that your counterpart possessed. Despite this, information gain assigns equal information value to r and $r*$.

Recall that $IG(r) = H(Q) - H(Q|r)$, where $H(Q)$ is the initial entropy of Q and $H(Q|r)$ the entropy of Q once you get r. The initial entropy of Q and $Q*$

is the same, for the priors are distributed in the same way. And the entropy of Q once you get r is the same as the entropy of $Q*$ once your counterpart gets $r*$, since, again, in both cases the probabilities are distributed in the same way. But this means that the information gain of r is the same as that of $r*$. The other measures avoid this problem, including the accuracy gain account. The expected accuracy gain of $r*$ is higher than that of r, so $r*$ comes out as more useful to your counterpart than r is to you on the accuracy gain model.[29]

3.4.2 Excluded Possibilities

Probability gain properly accounts for the value of surprise, and thus avoids the objection just raised to information gain. Nonetheless, it faces another serious problem. To illustrate, suppose that your research question Q this time has three possible answers, and that your prior probabilities are as follows: $P(q_1) = .5$, $P(q_2) = P(q_3) = .25$. Your counterpart is also interested in this question, and has equal priors as you. You get r, updating on which gives you the following posterior probabilities: $P_{new}(q_1) = .5, P_{new}(q_2) = .3, P_{new}(q_3) = .2$. Your counterpart gets $r*$, updating on which results in the exclusion of q_3, and the following posterior probabilities: $P'_{new}(q_1) = P'_{new}(q_2) = .5$. Intuitively, $r*$ is more useful to your counterpart than r is to you. Information $r*$ allows your counterpart to exclude one possible answer, whereas r only makes one slightly less likely. Yet probability gain gives the same value to both r and $r*$, namely 0, since neither answer raises the probability of the most likely hypothesis. The other measures avoid this problem, including the accuracy gain account. The expected accuracy gain of $r*$ is higher than that of r, so $r*$ comes out as more useful to your counterpart than r is to you on the accuracy gain model.[30]

[29] $EAG(r) = .5 \cdot (.36 - .25) \cdot 2 = .11$, whereas $EAG(r*) = .5 \cdot (.81 - .25) \cdot 2 = .56$.

[30] $EAG(r) = 0 + .3 \cdot (.5625 - .49) + .2 \cdot (.5625 - .64) = 0.00625$. By contrast, $EAG(r*) = 0 + .5 \cdot (.5625 - .25) + 0 = .15625$. It should be noted that according to the accuracy gain model, the *objective* information value of r and $r*$ is the same—i.e. 0—if q_1 is the true answer. Is this a bad result? Not

3.4.3 Belief Change

My arguments so far have proven particularly damaging for information gain and probability gain, but KL distance and impact come out as relatively untouched. Accuracy gain has the advantage of being able to incorporate additivity at the objective level, but other than that the three measures may look roughly on a par. I now want to argue that only accuracy gain is capable of providing a satisfactory explanation for why belief change in light of new evidence is useful. Insofar as KL distance and impact can at all provide an explanation, they will have to help themselves to the story given by accuracy gain.

Recall that both KL distance and impact measure the usefulness of a piece of information as a function of how much updating on that piece of information changes one's prior credences. The greater the change in credences, the more information value. But why should that be so? Surely, belief change is not of epistemic value *in itself*. The reason why belief change is useful must be that it's expected to get us closer to the truth about our research question. But the standard measure of closeness to truth is accuracy, and so we get the result that the reason why belief change is useful is that it's expected to bring about a gain in accuracy.[31] But why would we then want to use KL distance or impact as measures of information value, rather than expected accuracy gain itself? We may do so in experiments, if that comes handy, to the extent that KL distance and impact are good proxies for expected accuracy gain. But then it's

upon reflection. Consider how an external observer would judge the shift in probabilities from $P'(q_1) = .5$, $P'(q_2) = P'(q_3) = .25$, to $P'_{new}(q_1) = P'_{new}(q_2) = .5$. Such an external observer would think: the acquisition of $r*$ has two outcomes for the agent. On the one hand, it excludes a false hypothesis, q_3; but on the other, it lends more plausibility to another false hypothesis, q_2. The goodness of the good outcome is of strength comparable to the badness of the bad outcome, so the two outcomes cancel each other out—resulting in an objective information value of 0. The agent herself, on the contrary, cannot engage in this reasoning, for she doesn't know that q_2 is false: hence, she will also consider the possibility where $r*$ excludes a false hypothesis *and* lends plausibility to the true hypothesis. Overall, she should think that $r*$ has positive value, as predicted by EIV = EAG.

[31] Cf. Oddie (1997), who proves that acquiring new information which changes one's credences always improves expected cognitive value, where this is cashed out in terms of accuracy—though he doesn't specify any particular measure of accuracy, besides saying that it be strictly proper.

crucial to recognise that what grounds these measures' normative significance is accuracy gain, which is the only *finally* valuable property of information, from the epistemic point of view. So, while information which changes one's beliefs is useful, in accordance with KL distance and impact, it is so because it's expected to increase accuracy gain.

*

To sum up, the accuracy gain model is unique in being positioned to account for the value of surprise, of excluding possibilities, and of changing beliefs.

3.5 The Argument from Simplicity

The arguments of §3.3 and §3.4 lend support to the accuracy gain model on the grounds that it is consistent with more intuitions than its rivals and possesses greater explanatory power. In this section, I want to show that accuracy gain is also the simplest theory of information value available.

To see this, notice first that none of probability gain, information gain, KL distance, and impact provide an account of *objective* information value, alongside what we called *expected* information value. This is no flaw in its own right, especially once one bears in mind that these measures have largely been devised to learn facts about people's rationality, and what is rational for an agent to do is standardly assumed to depend entirely on the agent's own perspective. Yet as we examine each of these models from a purely theoretical point of view, asking which of them best accounts for the notion of information value, it makes sense to ask whether we shouldn't want them to have something to say about the objective perspective, too. We can ask: how useful should an agent S with research question Q and probabilities $P(\cdot)$ judge that piece of information r is

to her? And we can devise normative measures of *expected* information value to answer such a question. But we can also ask: how useful should an external observer acquainted with the truth about Q judge that piece of information r is to agent S? And we can equally devise normative measures of *objective* information value to answer such a question. The accuracy gain model offers an attractive package which addresses both questions: objective information value is correctly measured by accuracy gain, and expected information value is correctly measured by expected accuracy gain.

How do the other measures fare? The trouble is that there doesn't seem to be any real contender to accuracy gain at the level of objective information value. Notice a couple of desiderata for such an account of objective information value. First, we said earlier that misleading information should be seen as of negative value from the perspective of an external observer. And of course, information which increases the probability of the true answer should be seen as of positive value from such an objective perspective. But then, we end up with the result that any plausible measure of objective information value should assign negative value to information which moves the agent away from the truth, and positive value to information which gets the agent closer to the truth. And since the standard measure of closeness to the truth is accuracy, this pretty much already gets us to accuracy gain as the correct account of objective information value.[32]

[32]This becomes apparent if one does try to come up with potential models of objective information value inspired by models of expected information value alternative to expected accuracy gain. Consider first probability gain. What comes closest to an objective version of probability gain is $P(q_T|r) - P(q_T)$. But this equals accuracy gain when accuracy is defined as follows: $A(P(p), p) = -(1 - P(p))$. And this *is* indeed a way to measure accuracy, just not one which gets standardly used, for it's not strictly proper and it allows non-probabilistic credence distributions to be undominated (cf. Pettigrew (2016, pp. 34-35), Pettigrew (2019).). Notice, moreover, that taking the expectation of this measure doesn't actually deliver probability gain; so, such an objective measure only rather indirectly relates to probability gain. Impact doesn't fare better, for the objective value which it measures the expectation of is $.5 \cdot |P(q_T|r) - P(q_T)|$, which assigns positive value to information moving you away from the truth, and so is a non-starter. Information gain and KL distance give more interesting results. What comes closest to an objective analogue of information gain is $\log 1/P(q_T) - \log 1/P(q_T|r)$, which equals $\log(q_T|r)/P(q_T)$. As with probability gain, there is only an indirect connection here, for taking the expectation of this value doesn't give you information

Once this is established, expected accuracy gain clearly emerges as the simplest model of *expected* information value. The link between accuracy gain as a measure of objective information value and expected accuracy gain as a measure of expected information value is far more straightforwardly justified than is the link between accuracy gain as a measure of objective information value and any of probability gain, information gain, KL distance, and impact as measures of expected information value.[33] So, the accuracy gain model of objective and expected information value provides the simplest package available, and this reinforces the case in its favour.

3.6 Empirical Adequacy

Recall that in §3.1 I noted how cognitive psychologists disagree about the extent to which facts about how people gather information should constrain normative models of information value. Some think that a normative model should as far as possible aim to vindicate people's behaviour.[34] Others allow for more discrepancy between the descriptive and normative levels.[35] Most agree, however, that other things equal we'd rather have a normative model which didn't postulate massive irrationality on people's part. In this section, I show that the accuracy gain model promises to meet basic standards of empirical adequacy.

gain. But the expectation of this value coincides with *another* measure, namely KL distance! And $\log P(q_T|r)/P(q_T)$ is a special case of accuracy gain, when accuracy is measured thus: $A(P(q_T), q_T) = \log P(q_T)$. So if the logarithmic score is a legitimate accuracy measure, KL distance turns out to be a special case of the accuracy gain model, once coupled with the objective value which it is the expectation of (cf. Roche and Shogenji (2018)). Since the logarithmic score likely isn't *the* correct accuracy measure—a case which, as we've seen, has by contrast been made with respect to the Brier Score (Pettigrew (2016))—the comprehensiveness of the accuracy gain model remains preferable. However, if the logarithmic score turned out to be the only correct measure of accuracy, the contribution of this Chapter could be seen as a justification of the claim that, and an explanation why, KL distance is the correct measure of expected information value—see, e.g., Levinstein (2012) for an argument supporting the logarithmic scoring rule.

[33] As just noted, even proponents of KL distance face the additional, arduous step of showing that the logarithmic score is the uniquely legitimate accuracy measure if they are to justify KL distance's connection to accuracy gain.

[34] See, e.g., Oaksford and Chater (2007).

[35] See, e.g., Baron (1985, pp. 52-53).

First, I show that the model predicts people's standard ranking in covariation tasks. Then, I show how a common heuristic approximates expected accuracy gain, and can thus be reliably used, in proper circumstances, to ask the questions which maximise expected accuracy gain.

3.6.1 Covariation Assessment

In covariation tasks, people are asked to assess the degree to which two variables, X and Y, are related. They are given values for: (i) the frequency with which X and Y occur together (cell A), (ii) the frequency with which X is present but Y is absent (cell B), (iii) the frequency with which X is absent but Y is present (cell C), and (iv) the frequency with which both variables are absent (cell D).

variable X \ variable Y	present	absent
present	A	B
absent	C	D

Decades of research have established the result that participants do not consider the four pieces of information as equally important when determining the degree of covariation, although according to traditional descriptive statistics, like the φ coefficient, they are.[36] In particular, the standard ranking that people give, in terms of how important the information of each cell is, goes as follows: $A > B \approx C > D$.[37]

This has led some psychologists to hypothesise that participants do not

[36] $\varphi = (AD - BC)/[(A + B) \cdot (C + D) \cdot (A + C) \cdot (B + D)]^{1/2}$—where A, B, C, and D correspond to the respective cell values.

[37] For a representative sample of the literature, see Wasserman et al. (1990), Lipe (1990), Levin et al. (1993), McKenzie (1994).

approach the task as one of *description*, but rather of *inference*, whereby they use the information provided by the four cells as bearing on a question about which of two mutually exclusive hypotheses about a larger population of instances is true. And indeed it has been shown that once the task is approached this way, and the presence of X and Y is assumed to be rare in the larger population, while their absence common, joint presence does count as more informative than joint absence, vindicating the standard ranking that people give. Results of this kind have been proved using log diagnosticity, impact, and information gain as measures of information value.[38]

I now wish to give an illustration of the fact that the accuracy gain model also predicts people's standard ranking of the four cells' informativeness. To see this, assume that the two mutually exclusive hypotheses about the larger population that participants are trying to decide between are h_1, that X and Y are moderately correlated ($\rho = .5$), and h_2, that X and Y are independent ($\rho = 0$).[39] Further, $P(h_1) = P(h_2) = .5$.[40] We then need the *rarity* assumption that joint presence of X and Y is rare, and joint absence common. Say, for instance, that $P(X) = P(Y) = .1$. With these assumptions about the set-up in play, we are in a position to calculate the probabilities necessary to obtain the expected accuracy gain of each of the pieces of information contained within the cells. Such calculations are carried out in Appendix A. Reporting here the final step only, and using lower-case letters for the propositions corresponding to the information contained in each cell, we get:

- $EAG(a) = P(h_1|a)\cdot[A(P(h_1|a), h_1) - A(P(h_1), h_1)] + P(h_2|a)\cdot[A(P(h_2|a), h_2) - A(P(h_2), h_2)] = .846 \cdot (.2263) + .154 \cdot (-.4657) = .1197$

[38] McKenzie and Mikkelsen (2007) use log diagnosticity, Nickerson (1996) uses impact, and Oaksford and Chater (1994) use information gain.

[39] I borrow the case from McKenzie and Mikkelsen (2007). ρ is the Pearson's product-moment correlation coefficient, ranging between -1 and 1. To the extent that the coefficient is close to 1 (-1), there is a positive (negative) relation between the variables.

[40] So, we're also assuming the two hypotheses to be jointly exhaustive given the participants' knowledge.

- $EAG(b) = .333 \cdot (-.1856) + .666 \cdot (.1411) = .0318$
- $EAG(c) = EAG(b) = .0318$
- $EAG(d) = .514 \cdot (.0099) + .486 \cdot (-.0101) = .0001$.

So, $EAG(a) > EAG(b) = EAG(c) > EAG(d)$, validating participants' standard judgement that the information contained in cell A is more useful than that contained in cells B, C, and D, with the information of cells B and C itself being more useful than that of cell D.

3.6.2 The Feature Difference Heuristic

Some may think that computing expected accuracy gain is such a complex and demanding cognitive task that it's implausible to think that human beings collect information in order to maximise this value. While it may be true that people's cognitive systems are not equipped with the tools needed to calculate precise expected accuracy gains, it's possible that we are prone to making use of heuristics that track expected accuracy gain reliably enough. I now illustrate that this is indeed the case with respect to the well-documented *feature difference heuristic*.

The planet Vuma scenario introduced by Skov and Sherman (1986) nicely illustrates the use of this heuristic strategy. The experimental subjects are told to imagine that they live on planet Vuma, in which only two types of creature exist: gloms and fizos. Due to the planet's blinding light and blowing sand, it's impossible to distinguish these creatures by sight. The task is to categorize a randomly chosen creature as either a glom or a fizo, by asking whether it wears a hula hoop or drinks iced tea, knowing about the following probabilities: $P(glom) = P(fizo) = .50$; $P(hula|glom) = .90$; $P(hula|fizo) = .50$; $P(drink|glom) = .28$; and $P(drink|fizo) = .32$.[41] Which feature would it

[41] These precise values are taken from Nelson (2008)'s discussion of the planed Vuma scenario.

be most useful to query about?

One thing that Skov and Sherman, but also Slowiaczek et al. (1992), found is that many subjects in cases such as the planet Vuma scenario ask about those features with maximal absolute difference in feature probabilities, i.e. those features f such that $|P(f|h_1) - P(f|h_2)|$ is maximal. This is a property of features which is readily available to the subjects; for example, in the planet Vuma case it's evident that it's the hula property that should be chosen according to the feature difference strategy:

- $|P(hula|glom) - P(hula|fizo)| = .90 - .50 = .40$
- $|P(drink|glom) - P(drink|fizo)| = .32 - .28 = .04$

Moreover, asking about the hula property is also widely assumed to be the rational and more useful choice, given these probabilities. So it's plausible to think that such a heuristic could be operative in helping subjects make effective choices about which queries to pursue.

I now want to show that asking questions via use of the feature difference heuristic can be a good way to maximise expected accuracy gain. I will illustrate the point with the Vuma scenario. What we need is to get the values for $EAG(Hula)$ and $EAG(Drink)$, where upper-case initials indicate that we're talking about questions rather than answers. These are derived in Appendix B. Reporting here the final step only, we get:

$$EAG(Hula) = P(hula) \cdot EAG(hula) + P(\neg hula) \cdot EAG(\neg hula) =$$

$$= .7 \cdot (.02) + .3 \cdot (.11) = .047$$

And:

$$EAG(Drink) = .3 \cdot (.0009) + .7 \cdot (.0001) = .00034$$

So, $EAG(Hula)$ is much greater than $EAG(Drink)$. The Vuma example shows how the feature difference heuristic can be used to maximise the expected accuracy gain of questions.

*

No doubt much remains to be shown about the extent to which people's information-gathering choices respect the injunction to maximise expected accuracy gain, and about how using simple heuristics can be an effective way to respect such an injunction. But the purpose of this paper was to argue that accuracy gain is the correct *normative* model of information value—our question was what information value *is*, and not what people take it to be. For such a purpose it suffices to show that the model doesn't postulate massive irrationality on people's part, and I take this section to have established at least that.

3.7 Conclusion

What is the correct measure of the usefulness of acquiring certain information, or of asking certain questions in order to acquire certain information, relative to the aim of settling a given research question? The accuracy gain model stands out as the best candidate theory of information value: it is simpler, it fits more intuitions, and it has greater explanatory power than any of the other available theories. Information value is to be accounted for in terms of accuracy gain.

This provides a strong case for the *accuracy-based* component of the account of Chapter 2. It's clear that agents are successful to the extent that their evidence-gathering maximises information value, relative to their *research question*. And it's clear that they are rational to the extent that their evidence-gathering maximises expected information value, again relative to their research question.

What wasn't so clear, and stood in need of argument, was the claim that accuracy gain, as opposed to some other putative epistemic value, is required to account for the dimensions of success and rationality in the correct epistemic axiology and normativity of evidence-gathering. But we just saw that information value is to be accounted for in terms of accuracy gain: the principles of **Generalized Success** and **Generalized Rationality** put forward in §2.6 then follow, and the *evidential accuracy* theory of inquiry is vindicated.

3.8 Appendix A

In this appendix, I show how the values for $EAG(a)$, $EAG(b)$, $EAG(c)$, and $EAG(d)$ are calculated, starting from the assumptions laid out in §3.6.1. As noted in footnote 39, I borrow the case from McKenzie and Mikkelsen (2007)—the reader will find a more detailed exposition of the set-up there.

Consider first h_2. Given that X and Y are independent under h_2, $P(a|h_2) = P(X) \cdot P(Y) = .01$. And similarly, $P(b|h_2) = P(X) \cdot [1 - P(Y)] = .09$; $P(c|h_2) = [1 - P(X)] \cdot P(Y) = .09$; $P(d|h_2) = [1 - P(X)] \cdot [1 - P(Y)] = .81$. As for h_1, with $\rho = .5$, $P(a|h_1) = .055$; $P(b|h_1) = .045$; $P(c|h_1) = .045$; $P(d|h_1) = .855$.[42] By the law of total probability, we get $P(a) = .0325$; $P(b) = P(c) = .0675$; $P(d) = .8325$. By Bayes' theorem, $P(h_1|a) = .846$, $P(h_2|a) = .154$; $P(h_1|b) = P(h_1|c) = .333$, $P(h_2|b) = P(h_2|c) = .666$; $P(h_1|d) = .514$, $P(h_2|d) = .486$.

We are now in a position to calculate the expected accuracy gain of each of the pieces of information contained within the cells:

- $EAG(a) = P(h_1|a) \cdot [A(P(h_1|a), h_1) - A(P(h_1), h_1)] + P(h_2|a) \cdot [A(P(h_2|a), h_2) - A(P(h_2), h_2)] = .846 \cdot (.2263) + .154 \cdot (-.4657) = .1197$
- $EAG(b) = .333 \cdot (-.1856) + .666 \cdot (.1411) = .0318$

[42] See McKenzie and Mikkelsen's paper for more detail on the derivation of these values.

- $EAG(c) = EAG(b) = .0318$
- $EAG(d) = .514 \cdot (.0099) + .486 \cdot (-.0101) = .0001$.

3.9 Appendix B

In this appendix, I show how the values for $EAG(Hula)$ and $EAG(Drink)$ are calculated. As noted in footnote 41, the initial probabilities given in the main text are taken from Nelson (2008)'s discussion of the planet Vuma scenario—the reader will find a more detailed exposition of the set-up there.

We first have to derive the following probabilities from those given by the case of §3.6.2, which can be done by employing the law of total probability and Bayes' theorem: $P(glom|hula) = .64$, $P(glom|\neg hula) = .17$; $P(fizo|hula) = .36$, $P(fizo|\neg hula) = .83$; $P(hula) = .7$. Then, $EAG(hula) = .64 \cdot (.1204) + .36 \cdot (-.1596) = .02$. And $EAG(\neg hula) = .17 \cdot (-.4389) + .83 \cdot (.2211) = .11$. Knowing this, we can finally get:

$$EAG(Hula) = P(hula) \cdot EAG(hula) + P(\neg hula) \cdot EAG(\neg hula) =$$

$$= .7 \cdot (.02) + .3 \cdot (.11) = .047.$$

Next, we do the same for the $Drink$ question. $P(glom|drink) = .47$, $P(glom|\neg drink) = .51$; $P(fizo|drink) = .53$, $P(fizo|\neg drink) = .49$; $P(drink) = .3$. So, $EAG(drink) = .47 \cdot (-.0309) + .53 \cdot (.0291) = .0009$. And $EAG(\neg drink) = .51 \cdot (.0099) + .49 \cdot (-.0101) = .0001$. This gets us:

$$EAG(Drink) = .3 \cdot (.0009) + .7 \cdot (.0001) = .00034.$$

Chapter 4

A Comprehensive Theory of Inquiry

In this Chapter, I compare the *evidential accuracy* theory of inquiry developed so far with three main contenders: (i) zetetic veritism, according to which true belief is the aim of inquiry, (ii) the credal accuracy theory, according to which improving the accuracy of one's actual credences is the aim of inquiry, and (iii) zetetic gnosticism, according to which knowledge is the aim of inquiry. Such a comparison will bring out the distinctiveness of the account here defended, and it will also illustrate how the proposed account fares better than the three extant views in a number of respects. Towards the end of the chapter, I deal with three objections that may be levelled against my evidential accuracy account.

4.1 Two Desiderata

The evidential accuracy account satisfies two central desiderata for a satisfatory theory of inquiry. The first one is that it accommodate the fact that evaluative

judgements about how well inquiries go are *gradable* rather than binary. An inquiry can not only achieve or fail to achieve its aim. Crucially, it can achieve or fail to achieve its aim *to various degrees*. We can say that this inquiry went better than that one, but worse than yet another one. Call this the *gradability desideratum*. The second desideratum for a satisfactory theory of inquiry is that it accommodate the connection between the aim of inquiry and the value of information. In particular, such a theory should provide a plausible account of the value of information, which is derived from the theory's posited aim of inquiry as follows: information should count as valuable to the extent that acquiring it helps the agent achieve her inquiry's aim, as posited by the theory. Then, this account of information value is to prove competitive when compared to those offered by cognitive psychologists and statisticians. Call this the *information value desideratum*.

It is easy to see that the evidential accuracy theory of inquiry meets both desiderata. We just saw in Chapter 3 how the evidential accuracy theory provides an account of information value which is not only plausible, but actually *better* than the extant accounts. And it is clear that the evidential accuracy theory delivers evaluative judgements about how well inquiries go of a sufficient fineness of grain. Inquiry goes well (badly), according to the account, to the extent that the accuracy of the agent's evidential probability in the true answer to the given research question increases (decreases), after the agent acquires some evidence. Since there will be very many such *extents*, there will be very many degrees to which an inquiry can go well (badly). I now turn to look at how alternative proposals about the aim of inquiry fare with respect to these two desiderata.

4.2 Zetetic Veritism

According to what we can call *zetetic veritism*, outright true belief is the aim of inquiry.[1] As it stands, this view fares especially poorly in relation to our two desiderata. If the agent acquires the relevant true belief while inquiring, her inquiry goes well. If the agent doesn't acquire the relevant true belief, her inquiry goes badly. We are offered no tools for measuring degrees to which inquiries can go well or badly, for forming a true belief is an all-or-nothing matter. Thus zetetic veritism seems unable to meet the gradability desideratum. If this first desideratum isn't met, there is little hope for the second desideratum to be satisfied either, for the theory will lack the resources to account for different degrees to which information can be valuable, and so its measure of information value will be a non-starter when compared to the existing measures.

However, perhaps the zetetic veritist can couple the claim that true belief is the aim of inquiry with a claim about how to measure closeness to the aim of inquiry, namely closeness to true belief. If so, the veritist could say that an agent's inquiry goes well to the extent that the doxastic attitude that the agent holds with respect to the true answer to her research question at the end of her inquiry is close to true belief. The natural proposal is that closeness to true belief is to be measured by means of accuracy scoring rules. False belief and true belief would constitute the diametrically opposed endpoints of the "closeness to true belief" continuum, and in between one would have credences to be ranked according to their accuracy: the more accurate a given credence is, the closer to true belief it counts. This refined version of zetetic veritism shares much with the next proposal about the aim of inquiry that I will consider, the

[1]Zetetic veritism has been defended, for example, by Kvanvig (2003) and Lynch (2004), and will presumably be shared by those, like Goldman (2002), who think that true belief is the ultimate epistemic value. The label, however, is my own, though I follow Friedman (forthcoming b) in using the term "zetetic" for matters related to inquiry.

credal accuracy theory. I now turn to discuss this latter view, since the reasons why I think we should reject it speak just as well against this refined version of zetetic veritism.

4.3 The Credal Accuracy Theory

According to the credal accuracy theory, the aim of inquiry is to increase the accuracy of one's credence in the true answer to one's research question.[2] Inquiry goes well to the extent that the agent's posterior credence in the true answer to the research question is accurate. This theory clearly meets the gradability desideratum, for there are very many *extents* to which the agent's relevant posterior credence can be accurate.

However, it's not clear that any interesting notion of information value can be derived from such a theory. For notice that on this view, any piece of evidence can be anything from extremely useful to extremely harmful depending on what the agent makes of it—irrespectively of the agent's prior credence distribution. If the agent acquires a piece of misleading evidence and updates on it rationally, such a piece of evidence will count as of negative value according to the credal accuracy theory. But the credal accuracy theory will also judge the very same piece of evidence as of positive value if the very same agent with the very same prior credences irrationally misinterprets the evidential support relation and goes on to form a very accurate credence in the true answer to the research question as a result of acquiring that evidence. One can of course make the value of information so dependent on facts about the agent's belief formation, but this deprives such a notion of much epistemological interest.

To illustrate, consider the injunction to collect useful or valuable information

[2]As far as I know, no one has explicitly defended the credal accuracy theory in print. However, this view will presumably be shared by those, like Pettigrew (2016), who think that the accuracy of one's credences is "the *only* epistemic virtue" (p. 6, emphasis in original).

when inquiring: according to the credal accuracy theory, an agent who collected completely irrelevant evidence but somehow formed an accurate credence in the true answer to the research question as a result would be complying with such an injunction. The injunction seems very plausible. Yet it doesn't seem plausible that an agent could comply with it by acting in the way just described. So, the relevant notion of information value figuring in my formulation of the injunction can't be the one that falls out of the credal accuracy theory. Thus, the credal accuracy theory fails to meet the information value desideratum laid out in §4.1. *Mutatis mutandis*, the same point applies to the refined version of zetetic veritism considered above, for such a view only differs from the credal accuracy theory in positing true belief in the true answer, rather than maximally accurate credence in the true answer, as the doxastic attitude that makes an agent's inquiry go as well as it possibly can if held by the agent.[3]

These two theories stand in stark contrast with the evidential accuracy account that I proposed, according to which how well an inquiry goes is a function of what evidence the agent ends up acquiring in the process. On my view, what doxastic attitude the agent goes on to form as a result of acquiring that evidence is a matter pertaining entirely to the evaluative dimension of belief formation. In the next section, I will consider a view which is in one way similar to mine, in that it too evaluates inquiries by reference to the kind of evidence that the agent ends up acquiring. However, we will see that this view too faces considerable problems.

[3]The two theories would even coincide if it turned out that outright true belief just is credence 1—see, e.g., Greco (2015) for a version of this view.

4.4 Zetetic Gnosticism

According to what we can call *zetetic gnosticism*, knowledge is the aim of inquiry.[4] This view and the evidential accuracy account share the claim that an inquiry goes well only if the agent collects evidence which supports the true answer to the agent's research question. For in order to know p, your evidence must support p. As I stated the view, however, zetetic gnosticism also requires that the agent come to know the true answer by forming the corresponding belief on the basis of that evidence, if her inquiry is to go well. Thus, this version of zetetic gnosticism is in one way unlike the evidential accuracy account, in that it makes how well inquiries go partly dependent on the agent's belief formation. Another version of zetetic gnosticism holds that the aim of inquiry is to acquire evidence which puts one in a position to know the true answer. This version, like the evidential accuracy account, makes how well inquiries go independent of the agent's belief formation. My criticism of zetetic gnosticism applies to both versions of the theory; unless otherwise stated, I will leave aside subtle differences arising from the distinction between these two versions. I will spend considerable time discussing zetetic gnosticism, for this is the view that comes closest to my evidential accuracy theory.

First of all, notice that as it stands, zetetic gnosticism fares especially poorly in relation to our two desiderata, just like the simple version of zetetic veritism does. If the agent acquires knowledge of the true answer while inquiring, her inquiry goes well. If she doesn't, her inquiry goes badly. We are offered no tools for measuring degrees to which inquiries can go well or badly, for acquiring a certain piece of knowledge is an all-or-nothing matter. Thus zetetic gnosticism seems unable to meet the gradability desideratum. If this first desideratum

[4]Defenders of zetetic gnosticism include Williamson (2000), Millar (2011), and Kelp (2014), (ms). Again, the label is my own, though I follow Littlejohn (forthcoming) in calling the knowledge-first position "gnostic".

isn't met, there is little hope for the second desideratum to be satisfied either, for the theory will lack the resources to account for different degrees to which information can be valuable, and so its measure of information value will be a non-starter when compared to the existing measures.

4.4.1 Closeness to Knowledge: Accuracy of Evidence

However, perhaps the zetetic gnostic can couple the claim that knowledge is the aim of inquiry with a claim about how to measure closeness to the aim of inquiry, namely closeness to knowledge. If so, the gnostic could say that the agent's inquiry goes well to the extent that the agent's evidential state with respect to the true answer to her research question at the end of her inquiry is close to knowledge. Can we use a similar strategy as the one employed used above to refine zetetic veritism? This would involve ranking *evidential probabilities* for their accuracy, and placing them on a "closeness to knowledge" continuum whose endpoints would be knowledge on the one hand, and maximally inaccurate evidential probability on the other: the more accurate a given evidential probability is, the closer to knowledge it counts.[5] On this view, if the agent's inquiry provides the agent with evidence E which falls short of placing her in a position to know the true answer, but nonetheless supports it—say, the probability of the true answer on the agent's new total evidence is .8—then we will say that the agent's inquiry went reasonably well. Call this the *accuracy of evidence account* of closeness to knowledge.

This understanding of closeness to knowledge coupled with the version of zetetic gnosticism on which the aim of inquiry is *to be in a position to know* the true answer will give exactly the same predictions as our evidential accuracy

[5]In finite cases, where evidence can never give probability 1 to a falsehood, there may be no single most inaccurate evidential probability. Rather, there may be an infinite sequence of increasingly inaccurate evidential probabilities getting arbitrarily close to 1. If that were the case, there would actually be no end*point* on that end of the continuum. This makes talk of endpoints slightly misleading, but in a way that's harmless for our purposes.

account, at least in cases where the probability space is finite. For in such cases, one's total evidence E_{tot} will give probability 1 to a proposition p *iff* the agent is in a position to know p.[6] Infinitary cases may give rise to discrepancies between the two accounts; I will consider those in the final section of this Chapter. For now, the important thing to notice is that zetetic gnosticism so understood comes remarkably close to our evidential accuracy account (in a way analogous to how the credal accuracy theory comes close to the refined version of zetetic veritism).

However, Gettier cases and cases of merely statistical evidence pose problems for any attempt to align closeness to knowledge with the accuracy of one's evidential probabilities. According to this account, there will be no difference in closeness to knowledge between, say, a .8 Gettierized accurate justified credence and a .8 accurate justified credence based on non-Gettierizing evidence.[7] And there will be no difference between a .8 accurate justified credence based on merely statistical evidence, and a .8 accurate justified credence based on not-merely-statistical evidence, for example perceptual or testimonial evidence.

Consider first Gettier cases.[8] Suppose you're inquiring about what the time is. You go to the kitchen and see that the clock displays 3 pm. However, while it is in fact 3 pm, the clock is broken and displays that time no matter what. As a result, you don't know that it's 3 pm. So, the evidence you get falls short of placing you in a position to know that it's 3 pm. However, it will still justify being *confident* that it's 3 pm; say, for instance, that it justifies having a .8 credence in its being 3 pm. Now, presumably, not only do you fail to know that it's 3 pm, but your Gettierized credence, based on your evidential probability,

[6] At least according to the standard conception of evidential probability; cf. Williamson (2000, ch. 10).

[7] I switch to talking of justified credences rather than evidential probabilities, for the cases I am about to consider are more naturally characterised as involving an element of belief formation. But the reader should remember my assumption that the agent's evidential probability for p is x *iff* the the credence in p propositionally justified by the agent's evidence is x.

[8] Gettier (1963).

also counts as being far from knowledge, despite qualifying as high justified confidence in the true answer to your research question. Second, consider cases of merely statistical evidence. It is commonly thought, at least among gnostics, that an agent in a lottery case does not know that her ticket is a loser (before the results appear), no matter how many tickets there are in the lottery.[9] This is so despite the fact that the agent's evidence supports very high credence that her ticket is a loser. Assuming that the agent's ticket is in fact a loser, the agent's justified high credence that it is will also count as very accurate. Yet, intuitively, such an accurate justified credence is rather distant from knowledge, on the gnostic's picture.

By contrast, testimonial or perceptual evidence falling short of guaranteeing knowledge can still provide the agent with accurate evidential probabilities and justified credences which get pretty close to it. So, this understanding of closeness to knowledge fails to do justice to a number of central intuitions, and should for this reason be rejected. However, it is worth noticing that with the refutation of this account, the gnostic deprives herself of the extremely precise measuring tools offered by scoring rules, which give the evidential accuracy account of information value an upper hand. How then can the gnostic measure closeness to knowledge if not by looking at the accuracy of the agent's evidential probabilities? I now consider two alternative proposals.

4.4.2 Closeness to Knowledge: Excluded Possibilities

A second proposal that might be countenanced links closeness to knowledge to the exclusion of possibilities. To illustrate, suppose that you're trying to settle a question which has ten possible answers. In your inquiry, you get a piece of evidence which excludes two of them, in the sense that that piece of

[9]See, e.g., Williamson (2000, p. 249), Littlejohn (forthcoming).

evidence is inconsistent with such two answers. Your new total evidence leaves open the remaining eight answers to your question. Now, note that knowing the true answer to a question can plausibly be characterised as that doxastic attitude supported by evidence which excludes every other possible answer. If you know that q_1 is the true answer to Q, you know that, e.g., q_2 is *not* the true answer to Q, namely you know that q_2 is false. So your evidence excludes q_2. One might thus think that the more possible answers the evidence you collect excludes, the closer to knowledge you will be. Call this the *excluded possibilities account* of closeness to knowledge.

If we adopt this account of closeness to knowledge, we get the following theory of inquiry. Inquiry goes well to the extent that the agent gathers useful information. Information is useful to the extent that updating on it excludes possibilities, namely possible answers to the agent's research question. This is because obtaining a doxastic attitude based on evidence which excludes all answers but the true one—i.e. knowledge—is the aim of inquiry.

This account of closeness to knowledge avoids the problems faced by the accuracy of evidence account. Consider again the Gettier-style clock case. Prior to going to the kitchen, your evidence presumably already excludes some members of the partition induced by your research question *What time is it?*. Say that you can't exclude any time in the 2 pm-4 pm interval. Seeing that the clock displays 3 pm won't give you any evidence capable of excluding any of the times in the 2 pm-4 pm interval, for the evidence you get, namely that the clock *displays* that it's 3 pm, is perfectly consistent with its being, e.g., 2 pm. Similarly, merely statistical evidence in lottery cases doesn't exclude any members of the relevant partition: that the chance of your lottery ticket being a loser is very high is compatible with it being a winner.

The excluded possibilities account, however, fails to draw important dis-

tinctions that we'd want a theory of inquiry and information value to make. Consider as an illustration two pieces of evidence e_1 and e_2. Evidence e_1 excludes two out of the ten answers to your research question Q and doesn't speak in favour of any of the answers left open, so that the evidential probabilities are redistributed evenly. Evidence e_2, by contrast, not only excludes two of the false answers to Q, but additionally speaks strongly in favour of the true answer, despite not being able to exclude the remaining seven false answers. For example, we may assume that your new evidential probability in the true answer, upon updating on e_2, is .9. Now, surely e_2 is more useful or valuable than e_1. Moreover, it also seems that acquiring e_2 gets you closer to knowledge than acquiring e_1. So, there are two objections here: (i) the excluded possibilities account fails to line up with our intuitions about when information is valuable; (ii) the excluded possibilities account fails to line up with our intuitions about closeness to knowledge.

The limitations of this account are made apparent if we consider cases where the agent is inquiring into a polar question, whose only possible answers are p and $\neg p$. Here, the only useful information is that which puts you in a position to know the true answer, according to the excluded possibilities account, for only such information excludes one of the possible answers to your research question. Thus, the account fails to respect the gradability desideratum in the case of polar questions. We should look elsewhere if we hope to come up with a satisfactory, comprehensive gnostic theory of inquiry.

4.4.3 Closeness to Knowledge: The Bipartite Structure

The gnostic may at this point wish to come up with a measure of closeness to knowledge which takes into account *both* accuracy of evidence *and* proportion of excluded cells. Closeness to knowledge may be thought of along these lines

as a two-place function which takes as arguments (i) the accuracy of the agent's evidential probability, and (ii) the proportion of cells in the relevant partition that the agent's new total evidence excludes, and delivers a value in a specified range, where the greatest possible value will correspond to knowledge itself. Call this the *bipartite account* of closeness to knowledge.

The challenge for an account of this kind is to clarify how the first and second argument interact with one another to determine degrees of closeness to knowledge. One option would be to say that one of the two arguments is lexically superior to the other, in the sense that for any two agents, if one counts as closer to knowledge with respect to the lexically superior argument, then she counts as closer to knowledge overall, irrespective of how she fares with respect to the other argument. The other argument would only matter in cases of ties at the level of the lexically superior argument. But this option won't do, irrespectively of which argument is taken as lexically superior: if we choose argument (i) as lexically superior, then the account runs into problems similar to those faced by the accuracy of evidence account. On this version of the account, a .9 evidential probability in the true answer based on merely statistical evidence will be closer to knowledge than a .89 evidential probability in the true answer based on evidence which excludes most false answers. If we choose argument (ii) as lexically superior, then we get other implausible results. For example, consider again two pieces of evidence e_1 and e_2. Suppose e_1 excludes three of the ten answers to your research question Q and speaks in favour of one of the remaining *false* answers. And suppose e_2 excludes two of the ten answers to Q and speaks in favour of the *true* answer. Now, acquiring e_2 should count as getting you much closer to knowledge than acquiring e_1, for e_1 is effectively a piece of misleading evidence, while e_2 isn't. However, on this version of the bipartite account acquiring e_1 puts you closer to knowledge, and

thus is more useful, than acquiring e_2, for e_1 excludes more cells in the partition induced by Q than e_2.

So the only viable version of the bipartite account is one which allows for trade-offs between the accuracy of the agent's evidence and the proportion of cells that the agent's evidence excludes. The problem here is that it's not clear what sorts of considerations should favour choosing a certain trade-off over others. For example, is a very accurate evidential probability based on statistical evidence closer to knowledge than a .6 accurate evidential probability which excludes a higher proportion of cells? The gnostic may remain non-committal about such specific questions, only intending to lay out a general structure that the correct measure of closeness to knowledge should display. But this comes at the cost of losing much of the predictive power that we would want out of a comprehensive theory of inquiry. That choosing a specific trade-off over any other just seems arbitrary suggests that there may be no shared, clear and workable notion of closeness to knowledge available.

*

What this means is that zetetic gnosticism will face difficulties in satisfying our two desiderata. The gnostic will simply be unable to make comparative evaluations, in a large number of cases, as to how well inquiries go, for she won't be able to help herself to a clear notion of closeness to knowledge. And as a result, the corresponding notion of information value that the gnostic will put forward will fail to give any precise predictions, in a large number of cases, as to which of two pieces of information counts as more useful.

The gnostic may reply that we do have *some* clear intuitions about closeness to knowledge, such as that updating on misleading evidence pushes one further away from knowing than one previously was. Could the gnostic plug *at least*

those intuitions into her theory of inquiry to *partly* accommodate the gradability and information value desiderata? While this can be done, abductive standards for theory choice will still strongly favour the evidential accuracy account, at least with respect to the tasks of evaluating inquiries and offering an account of information value.[10] For notice just how much the evidential accuracy theory accomplishes. Independently of which accuracy scoring rule is chosen, such a theory offers an ordinal ranking of inquiries with respect to how well they go, and a related ordinal ranking of pieces of information with respect to how valuable they are. If one particular scoring rule is chosen, moreover, one gets *cardinal* rankings of inquiries and pieces of information. This is not just a simple, elegant and theoretically powerful outcome. More importantly, it helps with a number of purposes that we'd want our theory of inquiry to be able to serve. For example, a satisfatory theory of inquiry should be able to fit in with, and provide a plausible interpretation of, the results in cognitive psychology about the rationality of people's information-gathering practices discussed towards the end of Chapter 3. For that, a theory of inquiry needs to offer a value function representing the value of information, which is defined for each piece of information, relative to a given research question. But we can't get that with zetetic gnosticism, for the lack of a shared and clear notion of closeness to knowledge makes it impossible to devise such a precise value function for the value of information. By contrast, the evidential accuracy theory of inquiry offers precisely such a value function, and as we've seen in Chapter 3, promises to fit in with, and provide a plausible interpretation of, cognitive psychologists' results about the rationality of people's information-gathering practices.

[10]For a discussion of the application of abductive methodology to philosophical theorizing, see, e.g., Williamson (2016).

4.5 Objections

The evidential accuracy theory of inquiry thus compares favourably in relation to the three contenders that I have looked at: zetetic veritism, the credal accuracy theory, and zetetic gnosticism. To complete my defence of the evidential accuracy theory, I now turn to consider three objections that might be levelled against it.

4.5.1 Gettier Cases

Consider again Gettier cases. These have been mentioned above to pose problems for the accuracy of evidence account of closeness to knowledge. The thought was that agents in Gettier cases have high accurate evidential probability and high accurate justified credence in the true answer, but intuitively count as being distant from knowing. The gnostic may now accept this point, and reject the accuracy of evidence account of closeness to knowledge, but raise the following worry against the evidential accuracy theory of inquiry:

> Not only do Gettierized agents count as distant from knowing, they also intuitively count as distant from the aim of inquiry—whatever that happens to be. In other words, their inquiry does not intuitively count as going well. Yet such agents end up with high accurate evidential probabilities and high accurate justified credences. Hence, the evidential accuracy theory will classify their inquiries as going reasonably well.

Here is my response: I agree that something sub-optimal is going on in Gettier cases, but I question that this shows that the agents' inquiries in such cases go any worse than as predicted by the evidential accuracy account.

In particular, it is crucial to notice that Gettier cases are instances of *negative*

introspection (or *negative access*) failures. According to the principle of negative introspection, if p is not part of one's evidence, one has conclusive evidence that p is not part of one's evidence.[11] We can see how this principle fails in Gettier cases: in our example above, the agent's evidence *does not* include that it's 3 pm, but the agent lacks conclusive evidence that her evidence does not include that it's 3 pm. Indeed, her evidence supports that her evidence *includes* that it's 3 pm. Thus, the agent's evidential probability as to what the time is will be quite accurate, but her evidential probability as to what's part of her evidence will be misleading, i.e. her evidential probability will support the false proposition that her evidence includes that it's 3 pm. This rightly makes for a negative epistemic evaluation of the case: the agent doesn't fare very well *overall*, epistemically speaking. Moreover, as I've described the case so far, I've assumed that the agent conforms her credences to her evidential probabilities. But it's important to note that negative introspection failures *typically* result in failures to be rational: agents in Gettier cases will most likely *outright believe* that it's 3 pm, thus lending more confidence to its being 3 pm than their evidence warrants. If we're imagining the case this way, unwarranted outright belief rightly makes for an additional negative epistemic evaluation of the case: the agent doesn't fare very well *overall*, epistemically speaking. Further, agents in Gettier cases will also typically hold a number of false beliefs about their environment: in our case, the agent will presumably falsely believe that the clock is working. If we're imagining the case this way, false belief about one's surroundings rightly makes for an additional negative epistemic evaluation of the case: the agent doesn't fare very well *overall*, epistemically speaking.

By pointing to these various epistemic dimensions relative to which Gettierized agents fare poorly, I submit, we can explain away the judgement that their inquiries go worse than as predicted by the evidential accuracy account.

[11] I am using Salow (2019)'s formulation of the principle here.

Gettierized agents' inquiries go relatively well, but they are heavily misguided in a number of ways as to how and why it is that their inquiries go the way they do.

4.5.2 Merely Statistical Evidence

A similar worry might arise with respect to merely statistical evidence. Lottery cases have been mentioned above to pose problems for the accuracy of evidence account of closeness to knowledge. The thought was that agents in lottery cases have high accurate evidential probability and high accurate justified credence that their ticket is a loser (remember we assumed that their ticket is in fact a loser), but intuitively count as being distant from knowing, at least by the gnostic's own lights. The gnostic may now accept this point, and reject the accuracy of evidence account of closeness to knowledge, but raise the following worry against the evidential accuracy theory of inquiry:

> Not only do agents in lottery cases count as distant from knowing, they also intuitively count as distant from the aim of inquiry—whatever that happens to be. In other words, their inquiry does not intuitively count as going well. Yet such agents end up with high accurate evidential probabilities and high accurate justified credences. Hence, the evidential accuracy theory will classify their inquiries as going reasonably well.

My response here will be to bite the bullet and simply deny that merely statistical evidence doesn't make for successful inquiries, if non-misleading. If we follow the gnostic's thought that merely statistical evidence cannot put one in a position to know (e.g. a lottery proposition), we should conclude that such evidence will never support the true answer to one's question to degree 1, and so it will never

be conclusive on its own.[12] As a result, merely statistical evidence will never place one in the best possible position one can occupy with respect to one's inquiry—that of maximally accurate evidential probability in the true answer.

But it's not clear that if you update on some evidence E and improve the accuracy of your evidential probability in the true answer to degree x, how well your inquiry counts as going also partly depends on whether E is a piece of merely statistical evidence or not. To illustrate, consider the following pair of cases:

> *London Trip 1*. You want to find out whether your housemate went to London for the weekend. She told you she would like to go if she managed to finish her essay in time. On Saturday, you see a copy of her finished essay on the table, and notice that she is not at home. Further, her story on Instagram shows that she was on a bus this morning. Your evidential probability that she went to London is .9 (let us assume). Indeed, she went to London.

> *London Trip 2*. You want to find out whether your housemate went to London for the weekend. You have no idea whether she wanted to go. However, an oracle tells you that yesterday she participated in a lottery with ten tickets, buying one, and decided that she'd go to London for the weekend just in case she lost the lottery. Your evidential probability that she went to London is .9. Indeed, she went to London.

In London Trip 2 your evidence is merely statistical, in London Trip 1 it isn't. Yet to the extent that what you care about is finding out if your housemate went to London, it seems that you should be indifferent as to which of these two sets

[12] Again, at least as far as finite probability spaces are concerned. I will shortly consider infinitary settings.

of evidence you'd rather possess. This suggests that merely statistical evidence weighs just as much as any other kind of evidence in determining how well an agent's inquiry goes.

4.5.3 Infinitary Cases

A final objection to the evidential accuracy theory appeals to infinitary cases. Suppose that some coins have been tossed. Your research question is whether they all landed heads. You find out that the sequence of coin tosses was infinite. Your evidential probability that they didn't all land heads increases to 1, once you update on the information that the sequence was infinite.[13] Assume further that they didn't in fact all land heads. According to the evidential accuracy account, you have reached the best possible epistemic position you can occupy with respect to your inquiry into whether all coins in the sequence landed heads. However, the objection goes:

> You don't *know* that they didn't all land heads; you cannot rule out that they didn't; for all you know, they did. There's more evidence you can collect which would improve your epistemic position with respect to your research question. For example, if you found out that some coins landed tails, that would count as an improvement in your epistemic position, for you would then know the answer to your research question.[14]

I have two main lines of reply to this objection. The first one is to note that it's debatable whether or not you don't know that not all coins landed heads, upon finding out that the sequence was infinite. Above, I granted the point that you don't know of any *finite* lottery that your ticket is a loser (before the results

[13]This is because any other assignment of probabilities would result in a violation of the axioms of probability theory—see Williamson (2007).

[14]Cf. Friedman (2019, p. 308), Williamson (forthcoming).

appear). I didn't dwell on criticisms to this claim for my interlocutor was the gnostic—and that you don't know of any finite lottery that your ticket is a loser is by far the standard position among knowledge-firsters.[15] However, it is now worth noticing that this and related claims about chances and knowledge are far from universally accepted. For instance, a recent debate in epistemology is built on the assumption that there's a number n (say, 1000) such that you know of a sequence of n coin tosses that they didn't all land heads (if they didn't).[16] We need not take a stance on this issue here. If it turned out that one can know lottery propositions, provided that the lottery is sufficiently large and that the relevant proposition is true, then some of my discussion above would pan out differently, but in inconsequential ways. Rather, my point is simply to note that if there is disagreement with respect to finite sequences of coin tosses as to whether one can know that not all coins landed heads, then the claim that one doesn't know that they didn't all land heads in the infinite case (when they didn't) should look rather suspicious. Notice, in particular, that all that can be said in support of the claim that one always lacks knowledge in infinitary cases looks remarkably similar to usual sceptical challenges: "you just can't rule out the possibility, however far-fetched, that they all landed heads", "*how* do you know that they didn't?"; etc. If anything, we should intuitively be *much less* confident in most of our ordinary claims to knowledge than in the proposition that not all coins in the infinite sequence landed heads. Consider virtually any case of testimony, for example: the possibility that the speaker lied is certainly much less far-fetched than the possibility that all coins in an infinite sequence of tosses landed heads.

These considerations suggest the following two-pronged reply to the objec-

[15] E.g., Williamson (2000), Littlejohn (forthcoming).

[16] Or, at least, this debate is built on the assumption that *it's plausible to think* that there is such a number n: see Bacon (2014), Dorr et al. (2014), Rothschild and Spectre (2018). Even more on point, see Sosa (2015, p. 120) for a defence of the claim that we can know lottery propositions.

tion above. It's plausible to think that you *know*, upon acquiring the information that the sequence is infinite, that not all coins landed heads (first prong). This doesn't mean that there is no more evidence that you could collect on the matter, if you wanted to. Analogously, suppose an agent gains knowledge of the true answer to her research question by testimony. She has reached the best possible epistemic position with respect to her inquiry into her research question—maximally accurate evidential probability. This doesn't mean that there is no more evidence that she can collect on the matter, if she wanted to.

But let us even grant for the sake of argument that you *don't know*, upon acquiring the information that the sequence is infinite, that not all coins landed heads (second prong). Why go into any lengths to acquire such knowledge when, as just noted, we should intuitively be much less confident in most of the propositions we claim to know than in the proposition that not all coins in the infinite sequence landed heads? The objection can be turned on its head: if we don't know in infinitary cases, then this shows that knowledge doesn't matter if one already has maximally accurate evidential probability.

For my second reply to the objection I assume again for the sake of argument that one lacks knowledge in infinitary cases. This reply looks at the role that the evidence we possess plays in practical reasoning. Plausibly, the ultimate reason why we collect information is to make better informed decisions.[17] Now, under what conditions would knowledge that not all coins landed heads make a difference in one's decision-making—over and above accurate evidential probability 1 that they didn't? If we use standard expected utility theory, getting such knowledge on top of accurate evidential probability 1 would have

[17]The point made in the Introduction that the evaluative and normative dimensions intrinsic to evidence-gathering are epistemic and not prudential still stands, however. Compare: plausibly, the ultimate reason why we play chess is to have fun. This doesn't make the evaluative and normative dimensions intrinsic to chess prudential. Yet, if you didn't know the rules of chess, but knew that the ultimate reason we play chess is to have fun, you would expect such rules to be fun-conducive. By the same token, we should expect the *epistemic* aim of evidence-gathering to be conducive to better decision-making.

no impact whatsoever in our decision-making.

But perhaps it's question-begging to use standard expected utility theory in this context. Perhaps, for example, we should use dominance reasoning in choices involving infinite probability spaces.[18] If that were the case, knowledge that a 0 probability event fails to obtain could make a difference beyond maximally accurate evidential probability that such an event doesn't obtain. The way it could make such a difference is by excluding from the decision table a possibility which makes an act dominate another. Let me illustrate with the following example: you happen to enter a lottery with infinitely many tickets, and only one winning. You can either (i) get one ticket for free, or (ii) get two tickets for free. Assume that at least the second ticket is in fact a loser. If you were to come to know that the second ticket is a loser, that would change what's rational for you to do if dominance is the correct rule for practical reasoning in infinitary cases. For before you gain such knowledge, act (ii) dominates act (i), but after you get such knowledge, the two acts are on a par.[19] Since you had maximally accurate evidential probability that the second ticket is a loser all along, this would be a case where knowledge makes a practical difference, on top of accurate evidential probability 1. However, it can be argued that employing dominance reasoning in infinitary settings leads to paradoxical conclusions. Since this argument is quite complex, I develop it in Appendix C, at the end of this Chapter.

With the argument of the Appendix, one way in which knowledge might be thought to make a difference in decision-making, over and above the contribution made by maximally accurate evidential probability, is ruled out. The onus

[18]I have the following dominance principle in mind: for all options (i.e. acts) $o*$ and o, if $U(o*, w) \geq U(o, w)$ for all worlds $w \in W$ and $U(o*, w) > U(o, w)$ for some world(s) $w \in W$, then $o*$ is better than o for any agent with utility function U; any such agent would act irrationally if they chose o over $o*$ or remained indifferent between the two.

[19]This is because the states that we know not to obtain can be "left off" a decision table—see Weatherson (2012).

is on the gnostic to find cases in which it does. For now, it seems that an agent who conforms her credence to her maximally accurate evidential probability has done all she can epistemically do to act successfully. This may go some way towards responding to the objection. If you update on evidence which gives you accurate evidential probability 1, you have reached the best possible epistemic position with respect to your research question. Do well at the level of belief formation, and you will be set on a course to successful rational action, too.

4.6 Conclusion

This concludes my defence of the evidential accuracy theory of inquiry. The theory fares better than its three main contenders with respect to the gradability and information value desiderata. Moreover, it has the resources to meet some of the most salient objections. No doubt more could be said on both fronts, either to defend or attack the evidential accuracy theory. At least for now, however, the theory emerges as the best one on offer. In the next concluding Chapter, I will look at potential applications and extensions of the theory after briefly summarizing what has been accomplished so far.

4.7 Appendix C

In this Appendix, I give an argument for thinking that employing dominance reasoning in infinitary settings can lead to paradoxical results. The argument is based on the following case:

> A demon wakes you up at night and tells you the following true things: "An infinite sequence of coins has been tossed. We can make a deal: for some sum of money, you can select a set of coin

> tosses in the sequence—the more coin tosses in the set, the more you'll have to pay. Then, if all coin tosses in the set you select came up heads, I will steal your precious guitar. If at least one coin toss in the set came up tails, you will be able to retain your guitar. If you don't engage in this game, I will steal your guitar no matter what."

Now, suppose you value your guitar to such an extent that you'd be willing to pay up to £1000 for selecting the set which contains *all* coin tosses (assuming that in this case, monetary value lines up with utility). How much money should you be willing to pay for selecting an infinite set which contains all coin tosses after the n^{th}, for any positive integer n? Dominance reasoning suggests the following: you should be willing to pay strictly less than £1000, for the act of selecting all coin tosses after the n^{th} is dominated by the act of selecting all coin tosses.[20] Consider the possible scenario in which all coin tosses after the n^{th} come up heads, but some before the n^{th} come up tails. In such a scenario, selecting all coin tosses allows you to retain your guitar, but selecting only all coin tosses after the n^{th} results in the demon stealing it. Moreover, there is no possible scenario in which selecting all coin tosses after the n^{th} fares better than selecting all coin tosses. So, selecting all coin tosses dominates selecting only all coin tosses after the n^{th}. If you'd pay the same amount for each, you'd be valuing them equally (recall we assumed that monetary value lines up with utility in this case). But dominance reasoning tells you it's irrational to value them equally. So, you should be willing to pay strictly less than £1000 for selecting the set containing all coin tosses after the n^{th}.

But then, by the same dominance reasoning, you should be willing to pay even less for selecting the infinite set containing all coin tosses after, say, the $2 \cdot n^{th}$; and even less for the infinite set containing all coin tosses after the

[20] See Vallentyne (2000) for an endorsement of the view that it is rational to use dominance reasoning in infinitary settings.

$3 \cdot n^{th}$; and so on. Since the sequence is infinite, by successively engaging in this kind of dominance reasoning, you'll reach a point where you'll be willing to pay *no money at all* for an infinite sequence containing all coin tosses after the m^{th}, where m is a number much greater than n (exactly what number m is will depend on how much you were willing to pay for selecting all coin tosses after the n^{th}).[21] But this is patently irrational; it means that you'd rather have your guitar stolen—an object you value quite a bit—than pay *any* amount of money for carrying out an act—selecting all coin tosses after the m^{th}—which is virtually certain to save it from the demon's hands. The culprit is dominance reasoning, which enjoined you to value the set containing all coin tosses after the n^{th} strictly less than the set containing all coin tosses.

Can infinitesimal differences in value come to the rescue? Perhaps we can say that you should value the set containing all coin tosses after the n^{th} less than the set containing all coin tosses, but that the difference is infinitesimal. Say that you should value the former set £ $1000 - \epsilon$, where ϵ is an infinitesimal number. However, the following argument suggests that this fix won't do.[22] Call t_1 the time at which the first coin is tossed, and t_{n+1} the time at which the $n + 1^{th}$ coin is tossed. Suppose the demon tells you that a second infinite sequence of coins has been tossed, starting at t_{n+1}—call this infinite sequence of coin tosses *sequence 2*, and the first (original) infinite sequence of coin tosses, *sequence 1*. Your selection of a set of coin tosses can now be made from either

[21] This is because of the Archimedean property of real numbers. Suppose you were willing to pay up to £$1000 - r$ for selecting the set containing all coin tosses after the n^{th}. By the Archimedean property, there is an integer z such that $z \cdot r > 1000$. For some number greater than n—presumably, $2 \cdot n$—dominance reasoning will forbid you from paying more than £$1000 - 2 \cdot r$ for selecting the set containing all coin tosses after that number. Similarly, for some number *much* greater than n—presumably, $z \cdot n$— dominance reasoning will forbid you from paying more than £$1000 - z \cdot r$ for selecting the set containing all coin tosses after that number, which I called m in the main text. But we've just seen that $z \cdot r > 1000$. So dominance reasoning forbids you from paying any money at all for selecting all coin tosses after the m^{th}.

[22] The argument to follow is inspired by one in Williamson (2007). Williamson's argument has been criticized by Weintraub (2008). Footnote 22 below shows why even if Weintraub turned out to be right, that wouldn't be a problem for my overall project.

of the two sequences, and the same rules as above apply: if at least one coin in the set came up tails, you can retain your guitar; if not, the demon will steal it. How much money should you be willing to pay for the set containing all coin tosses from sequence 2? Clearly, £1000. We can map each coin toss of sequence 1 onto a corresponding coin toss of sequence 2: the first toss of sequence 1 onto the first toss of sequence 2, etc. Further, the conditions for employing dominance reasoning fail to obtain here: in some possible scenarios, selecting the set containing all coin tosses from sequence 2 results in a better outcome than selecting the set containing all coin tosses from sequence 1 (when some coins in sequence 2 land tails but none in sequence 1 does); in some other possible scenarios, selecting the set containing all coin tosses from sequence 1 results in a better outcome than selecting the set containing all coin tosses from sequence 2 (when some coins in sequence 1 land tails but none in sequence 2 does); and in most possible scenarios, the two acts result in equally good outcomes (when some coins land tails in both sequences). So you should value the two sets equally, i.e. £1000. But now notice that sequence 2 is in all relevant respects like the part of sequence 1 which begins after the n^{th} coin has been tossed: we can map the coin from sequence 1 tossed at t_{n+1} onto the (first) coin from sequence 2 tossed at t_{n+1}, etc. So, you should value the set containing all coin tosses in sequence 2 and the set containing all coin tosses in sequence 1 *after the* n^{th} equally. Since we just showed that you should value the set containing all coin tosses in sequence 2 £1000, you should value the set containing all coin tosses in sequence 1 after the n^{th} £1000. But this contradicts the proposal made at the start of this paragraph that you only value such a set £1000 $- \epsilon$.[23]

[23]Even if, contrary to what the argument just given suggests, infinitesimal value differences turned out to be able to rehabilitate dominance reasoning in infinitary settings, that wouldn't actually be bad news for the evidential theory of inquiry. For this would plausibly mean that infinitesimals can also be employed in probability theory: the probability that not all coins in an infinite sequence of tosses come up heads will be infinitesimally smaller than 1. The objection stemming from infinitary cases of §4.5.3 would lose its bite: one wouldn't have achieved yet maximally accurate evidential probability with respect to whether all coins landed heads if all

You should thus value the set containing all coin tosses in sequence 1 and the set containing all coin tosses in sequence 1 after the n^{th} equally: you should be willing to pay up to £1000 for either. Dominance reasoning breaks down in this kind of infinitary setting.

one knew was that the sequence is infinite. In the main text and here, I have assumed standard probability theory, but it's worth mentioning in this context that if the best theory of probability turned out to involve infinitesimal probabilities (see, e.g., Benci et al. (2016)), the evidential accuracy theory would only stand to gain.

Chapter 5

Conclusions

In this thesis, I have proposed and defended an account of the epistemic axiology and normativity surrounding the practice of evidence-gathering. Let me briefly summarise each of the account's building blocks. In the Introduction, I outlined ***The Setup***, intended to delineate our object of epistemological study: an agent is interested in settling a question, Q, and can perform one of a set of mutually exclusive and jointly exhaustive acts of evidence-gathering. I next presented, in Chapter 2, the *accuracy-based virtue-theoretic* account, which specified the conditions under which such an act of evidence-gathering counts as successful, rational, and skillful. At the core of the account is the notion of *accuracy gain*, which corresponds to the change (positive, negative, or null) in the accuracy of the agent's evidential probability in Q's true answer that updating on the evidence collected by performing the chosen act brings about. The greater the accuracy gain, the more successful the act will be. The greater the expected accuracy gain, the more rational the act will be. To the extent that the agent's act manifests a competence at collecting accuracy-maximising evidence, it will additionally count as skillful. Finally, if success and rationality

are achieved thanks to the agent's competence, the act will count as skillfully successful. Chapters 3 and 4 focused on defending the epistemological core of the accuracy-based virtue-theoretic account, which I called the *evidential accuracy theory of inquiry*. Chapter 3 argued that accuracy gain independently provides the best available account of *information value*. This result provides strong support in favour of the evidential accuracy theory, for it connects the success and rationality of an act of evidence-gathering with the usefulness of the information that the agent obtains by performing it. Chapter 4 looked at how the evidential accuracy theory fares *vis-à-vis* three main contenders, arguing that it better fits in with the two desiderata of gradability and information value and that it can meet three core objections. This concluded my defence of the evidential accuracy theory of inquiry.

However, it should be clear that my defence of such a theory is only to be seen as the starting point for epistemological work on the topic of evidence-gathering. For far from putting to rest philosophical questions about the nature of inquiry and evidence collection, the account here defended generates yet more avenues for thought, both by identifying additional problems in the epistemology of inquiry itself, and also by connecting the epistemology of inquiry to other philosophical areas. I now wish to conclude by looking at some of the avenues for future research suggested by the present account.

First, some research questions arise if one considers the limitations of the account here defended. One such limitation concerns the narrow focus with which I've been concerned. Recall that our agent in ***The Setup*** faces a single choice between a number of acts of evidence-gathering that she can decide to perform. However, suppose that we now add an additional element of complexity to the initial setup: the agent faces a *sequence* of choices between different acts of evidence-gathering that she can decide to perform. It may be

that attempting to maximise accuracy gain at each stage in the inquiry is not the best strategy for the agent to adopt in order to maximise accuracy gain overall. In general, the best strategy may be one that does not bear its fruits until much later in the process, but that once it does it delivers benefits so great as to make up for a middling start. Thus, in cases where the agent faces sequences of choices, she should take a big-picture view and choose the sequence of acts recommended by the best strategy, rather than assess each individual choice in isolation.[1] In the case of evidence-gathering, this will mean that our agent will act rationally to the extent that she chooses those strategies that are expected to maximise accuracy gain overall. These brief remarks should make it clear that the evidential accuracy account is capable of being extended from the single choice case to the much more complex case of sequences of choices. But much work remains to be done in formally spelling out the details of this extended approach.

Another limitation of the account as currently developed concerns the empirical question of whether agents choose information in accordance with the injunction to maximise expected accuracy gain. The final section of Chapter 3 tentatively suggested a positive answer, but more research needs to be carried out for the result to be made robust. Relatedly, if it turned out that agents choose sources of evidence in an attempt to maximise expected accuracy gain, how would this conclusion bear on the traditional view in psychology that people's thinking is riddled with irrationality?[2] Such an "irrationalist" interpretation is now widely thought to rely on obsolete and simplistic models of how rational agents think or act.[3] A very recent research project in epistemology attempts to use philosophically sophisticated models of rationality to

[1] Cf. Baron (1985, pp. 145-7).

[2] See, e.g., Lord et al. (1979), Kahneman et al. (1982), Gilovich (1991), Nickerson (1998), Kahneman (2011).

[3] See, e.g., Anderson (1990), Oaksford and Chater (2007), Gershman et al. (2015), and references in the next footnote.

vindicate people's judgements and behaviour against the irrationalist interpretation.[4] The final section of Chapter 3 again suggests that the present work fits in nicely with the aims and scope of this research project, but interesting work remains to be done to see just *how much* of people's information-gathering choices can be vindicated on the assumption that they act in order to maximise expected accuracy gain. For instance, what is the relation between confirmation bias and expected accuracy gain maximisation? Questions like this will have to be addressed in future work.

Other avenues for future research arise not so much from limitations of the current account, but rather from the connections between the epistemology of evidence-gathering and other philosophical areas that the account helps to shed light upon. One such connection is with the philosophy of language. According to the current account, asking a question can be viewed as an act of evidence-gathering whose possible outcomes correspond to the question's possible answers. Suppose that a group of agents is interested in settling a certain question, or a set of questions. The current account suggests that given the purposes of their investigation, there will be pressure on these agents to praise the posing of questions with high expected accuracy gain, and blame the posing of questions with low or null expected accuracy gain. By the same token, the account suggests that there will be incentives for these agents to praise assertions with useful informational content, and blame assertions with useless informational content. Thus, it seems that group inquiries can trigger the activation of norms on posing questions, and of norms on assertion stronger than the norms on assertion that are ordinarily thought to be in place (e.g. *assert p only if you know it*).[5] Interesting future research could result by looking at how much work these hypothesised norms can do in explaining what questions and

[4] See, e.g., Kelly (2008), Hedden (2019), Dorst (ms a), Dorst (ms b), Dorst and Mandelkern (ms).
[5] This latter norm on assertion is defended most notably by Williamson (2000, ch. 11).

assertions count as relevant in a given conversational context.[6]

A second connection to be explored in future work concerns the confirmation measures proposed in the philosophy of science literature. The current account has already established a surprising and welcome connection between accuracy measures and measures of information value. Various measures have additionally been proposed to account for the amount of *support* that a given piece of evidence provides for a given hypothesis; these get called confirmation measures.[7] Some have already established connections between confirmation measures and accuracy measures on the one hand, and confirmation measures and information value measures on the other.[8] It is thus tempting to put all these pieces together and explore the prospects for the current account to be able to provide a unifying treatment of accuracy measures, information value measures, and confirmation measures.

These are only some of the philosophical uses to which the proposed account, or some modified, extended version thereof, can be put. That the account uncovers so many new areas ripe for philosophical inquiry is itself an additional argument in its favour, at least according to the abductive methodology that I've assumed throughout this work. The field of epistemology has seen, in recent decades, a desirable expansion from its initial focus on individual belief formation that characterised so much of the post-Gettier literature to other topics just as worthy of attention. Central to this expansion are social epistemology, the theory of rational action, and the topic of epistemic norms. Inquiry has also been recently re-discovered as an object of proper epistemological theorizing, which has extended epistemology's reach once again. With this thesis, I hope to have contributed somewhat to the study of inquiry, and more generally, to

[6]See van Rooy (2004) for a different approach to this topic, one which makes use of information gain, as defined in Chapter 3, for measuring information value.

[7]See Fitelson (1999) for an overview of confirmation measures.

[8]See van Enk (2014) for the former connection and Crupi and Tentori (2014) for the latter.

the expansion of epistemology. The aim and hope was to expand epistemology's reach so far as to make it overlap with that of other disciplines, especially statistics, cognitive psychology, and the philosophy of science.

Bibliography

Anderson, J. R. (1990). *The Adaptive Character of Thought*. Hillsdale, NJ: Erlbaum.

Åqvist, L. (1984). Deontic logic. In Gabbay, D., editor, *Handbook of Philosophical Logic*, pages 605–714. Dordrecht: Reidel.

Bacon, A. (2014). Giving your knowledge half a chance. *Philosophical Studies*, (2):1–25.

Baron, J. (1985). *Rationality and Intelligence*. Cambridge: Cambridge University Press.

Baron, J., Beattie, J., and Hershey, J. C. (1988). Heuristics and biases in diagnostic reasoning : II. Congruence, information, and certainty. *Organizational Behavior and Human Decision Processes*, 42(1):88–110.

Beddor, B. and Pavese, C. (2018). Modal virtue epistemology. *Philosophy and Phenomenological Research*.

Belnap, N. D. and Steel, T. (1976). *The Logic of Questions and Answers*. New Haven; London: Yale University Press.

Benci, V., Horsten, L., and Wenmackers, S. (2016). Infinitesimal probabilities. *British Journal for the Philosophy of Science*, 69(2):509–552.

Björnsson, G. and Finlay, S. (2010). Metaethical contextualism defended. *Ethics*, 121(1):7–36.

Bona, G. D. and Staffel, J. (2017). Graded incoherence for accuracy-firsters. *Philosophy of Science*, 84(2):189–213.

Box, G. E. and Hill, W. J. (1967). Discrimination among mechanistic models. *Technometrics*, 9(1):57–71.

Bradley, D. (2020). Naturalness as a constraint on priors. *Mind*, 129(513):179–203.

Brandt, R. B. (1959). *Ethical Theory*. Englewood Cliffs, NJ: Prentice-Hall.

Brier, G. W. (1950). Verification of forecasts expressed in terms of probability. *Monthly Weather Review*, 78(1):1–3.

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

Buchak, L. (2014). Belief, credence, and norms. *Philosophical Studies*, 169(2):1–27.

Caie, M. (2013). Rational probabilistic incoherence. *Philosophical Review*, 122(4):527–575.

Campbell-Moore, C. and Salow, B. (forthcoming). Avoiding risk and avoiding evidence. *Australasian Journal of Philosophy*.

Carr, J. (2015). Chancy accuracy and imprecise credence. *Philosophical Perspectives*, 29(1):67–81.

Ciardelli, I., Groenendijk, J., and Roelofsen, F. (2018). *Inquisitive Semantics*. Oxford: Oxford University Press.

Cohen, S. (2016). Theorizing about the epistemic. *Inquiry*, 59(7-8):839–857.

Conee, E. and Feldman, R. (2004). *Evidentialism: Essays in Epistemology*. Oxford: Oxford University Press.

Cover, T. M. and Thomas, J. A. (2006). *Elements of Information Theory.* Hoboken, NJ: Wiley-Interscience.

Crupi, V. and Tentori, K. (2014). State of the field: Measuring information and confirmation. *Studies in History and Philosophy of Science Part A*, 47:81–90.

Davidson, D. (2005). *Truth, Language and History*. Oxford: Oxford University Press.

de Finetti, B. (1974). *Theory of Probability*. New York: John Wiley.

Dorr, C., Goodman, J., and Hawthorne, J. (2014). Knowing against the odds. *Philosophical Studies*, 170(2):277–287.

Dorst, K. (2019). Lockeans maximize expected accuracy. *Mind*, 128(509):175–211.

Dorst, K. (ms a). Overconfidence in overconfidence.

Dorst, K. (ms b). Rational polarization.

Dorst, K. and Mandelkern, M. (ms). Good guesses.

Easwaran, K. (2013). Expected accuracy supports conditionalization—and conglomerability and reflection. *Philosophy of Science*, 80(1):119–142.

Evans, J. and Over, D. E. (1996). Rationality in the selection task: Epistemic utility versus uncertainty reduction. *Psychological Review*, 103(2):356–363.

Ewing, A. C. (1948). *The Definition of Good*. London: Routledge Kegan Paul.

Fedorov, V. (1972). *Theory of Optimal Experiments Designs*. New York: Academic Press.

Feldman, R. (2000). The ethics of belief. *Philosophy and Phenomenological Research*, 60(3):667–695.

Fitelson, B. (1999). The plurality of bayesian measures of confirmation and the problem of measure sensitivity. *Philosophy of Science*, 66:S362–S378.

Flores, C. and Woodard, E. (ms). Epistemic vigilance: In defence of a norm on evidence-gathering.

Friedman, J. (2019). Inquiry and belief. *Noûs*, 53(2):296–315.

Friedman, J. (forthcoming a). Zetetic epistemology. In *Towards an Expansive Epistemology: Norms, Action, and the Social Sphere.* London: Routledge.

Friedman, J. (forthcoming b). The epistemic and the zetetic. *Philosophical Review*.

Gershman, S., Horvitz, E., and Tenenbaum, J. (2015). Computational rationality: A converging paradigm for intelligence in brains, minds, and machines. *Science*, 349(6245):273–278.

Gettier, E. (1963). Is justified true belief knowledge? *Analysis*, 23(6):121–123.

Gibbard, A. (2005). Truth and correct belief. *Philosophical Issues*, 15(1):338–350.

Gilovich, T. (1991). *How We Know What Isn't So. The Fallibility of Human Reason in Everyday Life*. New York: Free Press.

Goldman, A. I. (2002). *Pathways to Knowledge: Private and Public*. Oxford: Oxford University Press.

Good, I. J. (1950). Probability and the weighing of evidence. *Philosophy*, 26(97):163–164.

Good, I. J. (1966). On the principle of total evidence. *British Journal for the Philosophy of Science*, 17(4):319–321.

Good, I. J. (1975). Explicativity, corroboration, and the relative odds of hypotheses. *Synthese*, 30(1-2):39–73.

Good, I. J. (1983). *Good Thinking: The Foundations of Probability and its Applications*. Minneapolis: University of Minnesota Press.

Greaves, H. (2013). Epistemic decision theory. *Mind*, 122(488):915–952.

Greaves, H. and Wallace, D. (2006). Justifying conditionalization: Conditionalization maximizes expected epistemic utility. *Mind*, 115(459):607–632.

Greco, D. (2015). How I learned to stop worrying and love probability 1. *Philosophical Perspectives*, 29(1):179–201.

Greco, J. (2010). *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity*. Cambridge: Cambridge University Press.

Groenendijk, J. and Stokhof, M. (1984). *Studies on the Semantics of Questions and the Pragmatics of Answers*. PhD Thesis, University of Amsterdam.

Hamblin, C. L. (1958). Questions. *Australasian Journal of Philosophy*, 36(3):159–168.

Hedden, B. (2015a). Oxford: Oxford University Press.

Hedden, B. (2015b). A defense of objectivism about evidential support. *Canadian Journal of Philosophy*, 45(5):716–743.

Hedden, B. (2019). Hindsight bias is not a bias. *Analysis*, 79(1):43–52.

Holguín, B. (ms). Thinking, guessing, and believing.

Holton, R. (2014). Intention as a model for belief. In Vargas, M. and Yaffe, G., editors, *Rational and Social Agency: Essays on the Philosophy of Michael Bratman*. New York; Oxford: Oxford University Press.

Horgan, T. (2017). Troubles for bayesian formal epistemology. *Res Philosophica*, 94(2):233–255.

Horowitz, S. (2019). Accuracy and educated guesses. *Oxford Studies in Epistemology*, 6:85–114.

Huemer, M. (2009). Explanationist aid for the theory of inductive logic. *British Journal for the Philosophy of Science*, 60(2):345–375.

Jackson, F. (1986). A probabilistic approach to moral responsibility. *Studies in Logic and the Foundations of Mathematics*, 114:351–365.

Jackson, F. and Pargetter, R. (1986). Oughts, options, and actualism. *Philosophical Review*, 95(2):233–255.

Jeffrey, R. (1970). Dracula meets wolfman: Acceptance vs. partial belief. In Swain, M., editor, *Induction, Acceptance, and Rational Belief*, pages 157–185. Dordrecht: Reidel.

Joyce, J. (1998). A nonpragmatic vindication of probabilism. *Philosophy of Science*, 65(4):575–603.

Joyce, J. (2009). Accuracy and coherence: Prospects for an alethic epistemology

of partial belief. In Huber, F. and Schmidt-Petri, C., editors, *Degrees of Belief*, pages 263–297. Springer.

Kahneman, D. (2011). *Thinking, Fast and Slow*. New York: Farrar, Straus and Giroux.

Kahneman, D., Slovic, P., and Tversky, A. (1982). *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge: Cambridge University Press.

Kelly, T. (2008). Disagreement, dogmatism, and belief polarization. *Journal of Philosophy*, 105(10):611–633.

Kelp, C. (2014). Two for the knowledge goal of inquiry. *American Philosophical Quarterly*, 51(3):227–32.

Kelp, C. (ms). Theory of inquiry.

Klayman, J. and Ha, Y.-W. (1987). Confirmation, disconnrmation, and information in hypothesis testing. *Psychological Review*, 94:211–228.

Konek, J. and Levinstein, B. A. (2019). The foundations of epistemic decision theory. *Mind*, 128(509):69–107.

Kullback, S. and Leibler, R. A. (1951). On information and sufficiency. *Ann. Math. Statist.*, 22(1):79–86.

Kvanvig, J. L. (2003). *The Value of Knowledge and the Pursuit of Understanding*. Cambridge: Cambridge University Press.

Lasonen-Aarnio, M. (2010). Unreasonable knowledge. *Philosophical Perspectives*, 24(1):1–21.

Lasonen-Aarnio, M. (2020). Enkrasia or evidentialism? Learning to love mismatch. *Philosophical Studies*, 177(3):597–632.

Lasonen-Aarnio, M. (forthcoming). Perspectives and good dispositions. *Philosophy and Phenomenological Research*.

Leitgeb, H. and Pettigrew, R. (2010). An objective justification of bayesianism I: Measuring inaccuracy. *Philosophy of Science*, 77(2):201–235.

Levin, I., Wasserman, E., and Kao, S.-F. (1993). Multiple methods for examining biased information use in contingency judgments. *Organizational Behavior and Human Decision Processes*, 55(2):228–250.

Levinstein, B. A. (2012). Leitgeb and Pettigrew on accuracy and updating. *Philosophy of Science*, 79(3):413–424.

Lewis, D. K. (1973). *Counterfactuals*. Oxford: Basil Blackwell.

Lewis, D. K. (1988a). Relevant implication. *Theoria*, 54(3):161–174.

Lewis, D. K. (1988b). Statements partly about observation. *Philosophical Papers*, 17(1):1–31.

Lindley, D. V. (1956). On a measure of the information provided by an experiment. *Ann. Math. Statist.*, 27(4):986–1005.

Lipe, M. (1990). A lens model analysis of covariation research. *Journal of Behavioral Decision Making*, 3(1):47–59.

Littlejohn, C. (2013). No evidence is false. *Acta Analytica*, 28(2):145–159.

Littlejohn, C. (forthcoming). Truth, knowledge, and the standard of proof in criminal law. *Synthese*.

Lord, C. G., Ross, L., and Lepper, M. (1979). Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *Journal of Personality and Social Psychology*, 37(11):2098–2109.

Lord, E. (2018). *The Importance of Being Rational*. Oxford: Oxford University Press.

Lynch, M. P. (2004). *True to Life: Why Truth Matters*. Cambridge, MA: MIT Press.

McKenzie, C. (1994). The accuracy of intuitive judgment strategies: Covariation assessment and bayesian inference. *Cognitive Psychology*, 26(3):209–239.

McKenzie, C. and Mikkelsen, L. (2007). A bayesian view of covariation assessment. *Cognitive Psychology*, 54(1):33–61.

Millar, A. (2011). Why knowledge matters. *Aristotelian Society Supplementary Volume*, 85(1):63–81.

Miracchi, L. (2019). When evidence isn't enough: Suspension, evidentialism, and knowledge-first virtue epistemology. *Episteme*, 16(4):413–437.

Nelson, J. (2005). Finding useful questions: On bayesian diagnosticity, probability, impact, and information gain. *Psychological review*, 112(4):979–999.

Nelson, J. (2008). Towards a rational theory of human information acquisition. In Chater, N. and Oaksford, M., editors, *The Probabilistic Mind. Prospects for Bayesian Cognitive Science*. Oxford: Oxford University Press.

Nickerson, R. (1996). Hempel's paradox and Wason's selection task: Logical and psychological puzzles of confirmation. *Thinking and Reasoning*, 2(1):1–31.

Nickerson, R. S. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology*, 2(2):175–220.

Oaksford, M. and Chater, N. (1994). A rational analysis of the selection task as optimal data selection. *Psychological Review*, 101(4):608–631.

Oaksford, M. and Chater, N. (2007). *Bayesian Rationality: The Probabilistic Approach to Human Reasoning*. Oxford: Oxford University Press.

Oddie, G. (1997). Conditionalization, cogency, and cognitive value. *British Journal for the Philosophy of Science*, 48(4):533–541.

Parfit, D. (1984). *Reasons and Persons*. Oxford: Oxford University Press.

Parfit, D. (2011). *On What Matters.* Oxford: Oxford University Press.

Pavese, C. (2017). Know-how and gradability. *Philosophical Review*, 126(3):345–383.

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

Pettigrew, R. (2019). On the accuracy of group credences. *Oxford Studies in Epistemology*, 6:137–161.

Pettigrew, R. G. (2013). A new epistemic utility argument for the principal principle. *Episteme*, 10(1):19–35.

Ramsey, F. (1926). Truth and probability. In Eagle, A., editor, *Philosophy of Probability: Contemporary Readings*, (2011), pages 52–94. New York: Routledge.

Roche, W. and Shogenji, T. (2018). Information and inaccuracy. *British Journal for the Philosophy of Science*, 69(2):577–604.

Rorty, R. (1995). Is truth a goal of enquiry? Davidson vs. Wright. *Philosophical Quarterly*, 45(180):281–300.

Rothschild, D. and Spectre, L. (2018). A puzzle about knowing conditionals. *Noûs*, 52(2):473–478.

Salow, B. (2019). Elusive externalism. *Mind*, 128(510):397–427.

Savage, L. J. (1954). *The Foundations of Statistics*. New York: John Wiley.

Schoenfield, M. (2018). An accuracy based approach to higher-order evidence. *Philosophy and Phenomenological Research*, 96(3):690–715.

Shannon, C. E. (1948). A mathematical theory of communication. *The Bell System Technical Journal*, 27(3):379–423.

Skov, R. and Sherman, S. (1986). Information-gathering processes: Diagnosticity, hypothesis-confirmatory strategies, and perceived hypothesis confirmation. *Journal of Experimental Social Psychology*, 22(2):93–121.

Slowiaczek, L., Klayman, J., Sherman, S., and Skov, R. (1992). Information selection and use in hypothesis testing: What is a good question, and what is a good answer? *Memory cognition*, 20(4):392–405.

Sosa, E. (2007). *A Virtue Epistemology.* Oxford: Oxford University Press.

Sosa, E. (2015). *Judgment & Agency*. Oxford: Oxford University Press.

Staffel, J. (2017). Accuracy for believers. *Episteme*, 14(1):39–48.

Turri, J. (2011). Manifest failure: The Gettier problem solved. *Philosophers' Imprint*, 11(8):1–11.

Vallentyne, P. (2000). Standard decision theory corrected: Assessing options when probability is infinitely and uniformly spread. *Synthese*, 122(3):261–290.

van Enk, S. J. (2014). Bayesian measures of confirmation from scoring rules. *Philosophy of Science*, 81(1):101–113.

van Rooy, R. (2004). Utility, informativity and protocols. *Journal of Philosophical Logic*, 33(4):389–419.

Wason, P. C. (1968). Reasoning about a rule. *Quarterly Journal of Experimental Psychology*, 20(3):273–281.

Wasserman, E., Dorner, W., and Kao, S. (1990). Contributions of specific cell information to judgments of interevent contingency. *Journal of experimental psychology. Learning, memory, and cognition*, 16(3):509–21.

Weatherson, B. (2012). Knowledge, bets, and interests. In Brown, J. and Gerken, M., editors, *Knowledge Ascriptions*, pages 75–103. Oxford: Oxford University Press.

Wedgwood, R. (2012). Outright belief. *Dialectica*, 66(3):309–329.

Wedgwood, R. (2017a). Must rational intentions maximize utility? *Philosophical Explorations*, 20(sup2):73–92.

Wedgwood, R. (2017b). *The Value of Rationality*. Oxford: Oxford University Press.

Weintraub, R. (2008). How probable is an infinite sequence of heads? A reply to Williamson. *Analysis*, 68(299):247–250.

Wells, G. and Lindsay, R. (1980). On estimating the diagnosticity of eyewitness nonidentification. *Psychological Bulletin*, 88(3):776–784.

Williamson, T. (2000). *Knowledge and its Limits*. Oxford: Oxford University Press.

Williamson, T. (2007). How probable is an infinite sequence of heads? *Analysis*, 67(3):173–180.

Williamson, T. (2016). Abductive philosophy. *Philosophical Forum*, 47(3-4):263–280.

Williamson, T. (2017). Ambiguous rationality. *Episteme*, 14(3):263–274.

Williamson, T. (forthcoming). Knowledge, credence, and the strength of belief. In Flowerree, A. and Reed, B., editors, *Towards an Expansive Epistemology: Norms, Actions, and the Social World*. London: Routledge.

Worsnip, A. (2019). The obligation to diversify one's sources: Against epistemic partisanship in the consumption of news media. In Fox, C. and Saunders, J., editors, *Media Ethics: Free Speech and the Requirements of Democracy*, pages 240–264. London: Routledge.

Yablo, S. (2014). *Aboutness*. Princeton, NJ: Princeton University Press.

Yalcin, S. (2018). Belief as question-sensitive. *Philosophy and Phenomenological Research*, 97(1):23–47.