

Norms of inquiry

David Thorstad

Forthcoming in *Philosophical Topics*.

Please cite published version

Abstract

Epistemologists have recently proposed a number of norms governing rational inquiry. My aim in this paper is to unify and explain recently proposed norms of inquiry by developing a general account of the conditions under which inquiries are rational, analogous to theories such as evidentialism and reliabilism for rational belief. I begin with a reason-responsiveness conception of rationality as responding correctly to possessed normative reasons. I extend this account with a series of claims about the normative reasons for inquiry that we possess. I apply the account to shed light on two classes of norms that have featured prominently in recent discussions: norms of clutter avoidance forbidding agents from engaging in trivial inquiries, and norms of logical non-omniscience governing properties such as the deductive closure and consistency of an agent's belief state. I conclude with a discussion of the sense in which norms of inquiry should be regarded as epistemic norms.

1 Introduction

Inquiry, like belief, is governed by norms. Epistemologists have recently proposed a number of norms governing rational inquiry. Rational inquirers seek significant truths (Treanor 2014, 2018) and avoid trivial inquiries (Harman 1986; Whitcomb 2018); suspend judgment (Friedman 2017); gather evidence (Hall and Johnson 1998; Woodard and Flores forthcoming); avoid incessant double-checking (Friedman 2019; Woodard forthcoming); and appropriately allocate attention (Siegel 2017). One fruitful project is to add new norms to this list.¹ But as the list becomes increasingly populous, we should also turn our

¹For example, we might ask what, if anything is the aim of inquiry (Friedman forthcoming; Falbo forthcoming) or how inquiry relates to knowledge (Kelp 2021a; Sapir and van Elswyk 2021; Wilard-Kyle

sights higher and ask for a general account of norms of inquiry analogous to theories such as evidentialism and reliabilism for rational belief. Under what conditions are inquiries rational?

There are at least four benefits to the general study of rational inquiry. First, we aim not merely to list norms of inquiry but also to understand and explain why these norms are what they are. Second, understanding the grounds of norms clarifies their contents. For example, we know that rational inquirers sometimes gather evidence, but inquirers cannot gather all available evidence. Which subjects should inquirers gather evidence about, and how much evidence should they gather? Third, we need to adjudicate conflicts between competing norms of inquiry. For example, it is controversial whether and in what sense inquirers are required to suspend judgment (McGrath 2021). Fourth, a number of recent authors have proposed that norms of inquiry have a significant practical component (Friedman 2020; Harman 2004; Lord 2020; Thorstad 2022a). One of my aims in this paper is to show how we can make space for this claim without falsifying traditional epistemological theories.

The first order of business is to develop an account of rational inquiry. In Sections 2-3, I develop and motivate an account of rational inquiry. Sections 4-5 show how that account can be used to shed light on two puzzle cases for recent discussions of rational inquiry: norms of clutter avoidance and requirements of logical non-omniscience. A common objection to this account is that it answers the wrong question because it does not seek to characterize a distinctively epistemic type of rationality governing inquiry. I consider this objection in Section 6. Section 7 concludes.

2 Reason-responsiveness

What is the relationship between rationality and reasons? Suppose you are sitting in the living room reading an informative philosophical book when you notice smoke billowing

forthcoming).

out of the kitchen. If you are rational, then you will put down the book and investigate the smoke. A plausible explanation for that fact is that rationality requires acting in accordance with the balance of reasons that you possess. Because you have most reason to investigate the smoking kitchen, that is what you are required to do.

But rational inquirers do not only act in accordance with reasons. Suppose you were to investigate the smoking kitchen because you believed that you would be rewarded for doing so by a genie. Then we would like to say that your inquiry was irrational, even though you performed the inquiry you had most reason to perform. A plausible explanation is that rationality requires not only acting in accordance with reasons, but also responding correctly to those reasons in acting. Although your evidence-gathering was supported by strong safety-related reasons, in inquiry you did not respond to these reasons but rather to the imagined promise of a genie. As a result, your inquiry was irrational.

These thoughts are often taken to motivate a reason-responsiveness conception of rationality (Lord 2018; Kieseewetter 2017).

(Rationality as Reason-Responsiveness) For all agents S , times t and acts X ,
 S 's X -ing at t is rational if and only if in X -ing at t , S responds correctly to the
normative reasons that she possesses at t .

Here I have stated the reason-responsiveness conception for the special case of rational action, since inquiry is an activity. But the reason-responsiveness conception is also meant to apply to attitudes such as belief. Although I will not defend a reason-responsiveness conception of rational belief here, it can be helpful to think through the case of rational belief to get a handle on the account.

In the case of rational belief, the requirement of responding correctly to normative reasons is familiar to epistemologists as the epistemic basing relation. Believing rationally requires holding the beliefs supported by the balance of possessed reasons, namely our total evidence, and holding these beliefs based on the evidence.² Leading theories of the

²This discussion can be repeated, if desired, using pragmatist or other non-evidentialist views of rational

reason-responsiveness relation often mirror theories of epistemic basing. For example, we can require that *S*'s reasons play an appropriate causal role in producing or sustaining her *X*-ing (Arpaly and Schroeder 2015), or that *S* *X*'s as the result of a specified competence (Lord 2018).

Suppose we take on board a reason-responsiveness conception of rational inquiry. What follows? At least two features of the reason-responsiveness view deserve emphasis.

First, rationality is deontic. On a reason-responsiveness view, what we are rationally required to do and what we ought to do are the same thing (Lord 2017). More precisely:

(Rationality is Deontic) For all agents *S*, times *t* and acts *X*, *S* is rationally required at *t* to *X* just in case *S* ought, at *t*, to *X*.

Roughly put, the claim that rationality is deontic follows on a reason-responsiveness view because what *S* ought to do and what *S* is rationally required to do are determined in the same way: by the balance of her possessed normative reasons.³ Because you have more reason to investigate the billowing smoke than to read your book, you ought to investigate the smoking kitchen and that is also what you are rationally required to do.

In this way, the reason-responsiveness view explains why epistemologists often pass between claims about what agents are rationally required to do and what they should or ought to do. On the reason-responsiveness view, this slide is not an equivocation or a mistake since rational requirements and oughts coincide. At the same time, the reason-responsiveness view explains why rightness and rationality remain importantly distinct normative statuses. Rationality is a strictly more demanding status than rightness, since acting rationally requires not only doing what we ought, but also acting in response to the reasons for which we ought to act.

A second consequence of the reason-responsiveness view is that rationality is all-things-considered (Thorstad 2022a). I discuss this consequence at length in Section 6,

belief.

³On some views, establishing the implication that Rationality is Deontic takes a bit more work than this. See (Lord 2017) for a full argument.

but here is a first pass statement and motivation. On the reason-responsiveness view, all types of normative reasons combine to determine what agents are rationally required to do. We may, in addition, try to specify restricted types of moral, prudential or epistemic rationality which consist in responding correctly to specific classes of reasons.⁴ But in general we need an all-things-considered view of rationality to capture many roles that a theory of rationality is meant to play. For example, it is often held that rational agents defer to what is rational in their present situation (Greco and Hedden 2016; Levinstein 2017) and plan to do what would be rational in future contingencies (Gibbard 2003; Horowitz and Dogramaci 2016). On most plausible ways of spelling out these planning and deference requirements, each requirement tracks an all-things-considered conception of rationality. If agents have most epistemic reason to read a book but most all-things-considered reason to investigate the smoking kitchen, then they defer by investigating the smoke. Those of us reading about their plight plan to do the same when faced with similar circumstances.

The claims that rationality is deontic and all-things-considered help the reason-responsiveness conception to meet Kolodny's challenge: why be rational (Kolodny 2005)? The reason-responsiveness view answers Kolodny's challenge in the natural way. In doing what is rational, we do what we ought and also what we have most reason to do. By contrast, if we deny that rationality is deontic or all-things-considered, then there will be agents who ought or have most reason to one thing, but are rationally required to do another. In this case, Kolodny's challenge sharpens to the following: why do what is rational, instead of what we ought or have most reason to do? I am not aware of any compelling answers to this sharpening of Kolodny's challenge. Unless an answer can be produced, Kolodny's challenge will tell in favor of the view that rationality is deontic and all-things-considered.⁵

⁴I discuss this view in Section 6.

⁵One interesting strategy would be to accept Kolodny's challenge as a datum about rationality, but to turn our sights from questions about rationality to related questions: for example, what is the aim of inquiry, or what constitutive standards govern inquiry? If we read authors such as Christoph Kelp (Kelp 2021a,b) as having this more restricted project in mind, that would allow them to express correct views that are perfectly consistent with, even complementary to the present project.

In this section, we have developed a reason-responsiveness conception on which rationality involves responding correctly to normative reasons. Two noteworthy consequences followed: rationality is deontic and all-things-considered. To unpack the reason-responsiveness conception further, we will need to say something specific about normative reasons.

3 Reasons for inquiry

In the rest of this paper, I leave behind the more complicated question of when a token act is rational and take up the simpler questions of which acts we are rationally permitted or required to token. This will allow me to remain neutral between competing accounts of the reason-responsiveness relation in order to focus on the questions about rational permissions and requirements that have taken center-stage in recent discussions of rational inquiry.

On a natural reading of Rationality as Reason-Responsiveness, agents are rationally permitted to X just in case X -ing could be a correct response to the normative reasons that they possess. This in turn plausibly occurs when agents have sufficient reason to X , giving:

(Rational Permissions) S is rationally permitted at t to X if and only if S has sufficient possessed reason at t to X .⁶

Agents are rationally required to X just in case X is the only thing that they are permitted to do. If we accept that agents have decisive reason to X just in case X -ing is the unique thing they have sufficient reason to do, we recover:

(Rational Requirements) S is rationally required at t to X if and only if S has decisive possessed reason at t to X .

⁶Here and throughout the remainder of the paper I suppress quantifiers in the statements of principles when their domains are clear from previous statements.

But this is as far as we can proceed without controversy. To go any further, we need to say something specific about the nature of normative reasons. I will lay my cards on the table as precisely as I can.⁷ Readers with differing views are invited to trace the consequences of alternative assumptions throughout the rest of our discussion.⁸

Like many epistemologists, I adopt a teleological approach on which reasons track what is of value. A minimal consequence of that approach is the following:

(Weak Teleological Claim) *S* has sufficient possessed reason at *t* to *X* if and only if *X*-ing is at least as good as any alternative available to *S* at *t*.

It follows that *S* has decisive possessed reason to *X* just in case *X*-ing is best.

The Weak Teleological Claim is often associated with consequentialism, but in fact it is quite a bit weaker than that. Many nonconsequentialist views accept the Weak Teleological Claim. For example, deontologists can agree that we always have most reason to do what is best (Sylvan 2020). On a first-pass deontological view, we have most reason not to lie even if lying would prevent several lies from being told, because in lying we would dishonor the truth. For this reason, it is also better not to lie in this case. By contrast, consequentialists and other teleologists think on a first-pass that we would have most reason to lie in this case, because lying would promote the most truth-telling.

The question at issue here is whether reasons are always reasons to promote value, or whether there can also be reasons to instantiate, honor, or otherwise relate ourselves to value. In this debate I side with the consequentialist. Perhaps reasons for belief and other attitudes are reasons to instantiate values such as accuracy. But inquiry is an activity, not a state, and we act on the world in order to promote valuable ends:

(Promotion) All reasons for action are reasons to promote value.

⁷For a full defense see Thorstad (forthcoming b).

⁸For example, we might explore nonconsequentialist views on which value is not always to be promoted; adopt non-maximizing criteria of correctness; or move beyond the space of value-based epistemologies altogether.

Putting together Promotion with Rational Permissions, Rational Requirements, and the Weak Teleological Claim yields a reason-responsiveness, consequentialist view:⁹

(Reason-Responsiveness Consequentialist View) For all agents *S*, times *t* and acts *X*:

(Evaluative Permission) *S* is rationally permitted at *t* to *X* if and only if *X*-ing is at least as good as any alternative available to *S* at *t*.

(Evaluative Requirements) *S* is rationally required at *t* to *X* if and only if *X*-ing is the best alternative available to *S* at *t*.

(Promotion) All reasons for action are reasons to promote value.

What are the prospects for this account?

It may seem surprising to adopt a consequentialist account of rational inquiry, but I think there are good reasons for optimism.¹⁰ Consequentialist approaches to rational belief already enjoy considerable popularity (Ahlstrom-Vij and Dunn 2018). We should expect a consequentialist account of rational inquiry to perform better still, insofar as consequentialists have always been most at home in accounting for actions rather than attitudes. Opponents of epistemic consequentialism may even see a consequentialist account of rational inquiry as a natural place to relocate tradeoffs and other consequentialist mainstays which they see as unattractive in a theory of rational belief (Berker 2013a,b; Sylvan 2020). Moreover, global consequentialists have been reasonably successful in extending consequentialism beyond a theory of right action to account for new objects and types of normative assessment, including the rationality of decisionmaking procedures.¹¹

⁹Strictly speaking, the view also includes Rationality as Reason-Responsiveness, but as I have said I am concerned here only with what the view implies for rational permissions and requirements.

¹⁰Another traditional advantage of consequentialist approaches is that they account for the existence and proper resolution of tradeoffs among different dimensions of value (Singer 2019). Although there may be some controversy about whether tradeoffs should be admitted into theories of rational belief (Berker 2013a,b), there is much less controversy over the existence and importance of tradeoffs in the theory of rational action.

¹¹On global consequentialism, see Driver (2001); Kagan (2000); Parfit (1984); Pettit and Smith (2000).

Extending global consequentialism to the neighboring topic of rational inquiry provides a natural test case for the global consequentialist program.

If a consequentialist account of rational inquiry is to be made to work, it will need to have two features on top of those introduced in Section 2. First, we need to adopt a *rich axiology* on which not only practical goods such as welfare and desire-satisfaction, but also intellectual achievements such as truth and knowledge bear ultimate value. Consequentialism does not make the pragmatist claim that intellectual achievements are valuable only insofar as they conduce to practical goods. By opening space for intellectual achievements to bear value, the consequentialist puts herself in a better position to recover truth-directed epistemic norms.

Second, we need to draw a strict *level separation* between questions about belief and inquiry (Thorstad 2021, 2022b). It is one question what we are rationally permitted to believe and another question entirely how we are permitted to inquire. For example, it may be wastefully irrational to count the blades of grass on our lawns even if we could form a rational belief about the number of blades as a result of this inquiry. The present account is only a theory of rational inquiry, not an account of rational belief. Without level separation, a consequentialist theory of inquiry would imply the unpopular result that we have strong and widespread non-intellectual reasons for belief. While I do not want to foreclose on that possibility, it should not be forced on us merely by recognition of the fact that inquiry is an activity and, like most activities, inquiry answers partly to non-intellectual reasons.

Level separation is familiar to consequentialists in the practical domain as the distinction between the rationality of decision procedures and the rationality of the intentions that result. Here I propose extending that same separation to the case of theoretical inquiry and belief. This level separation is also a familiar part of traditional arguments for evidentialism. Level separation is used to hold that apparent non-epistemic reasons for belief, such as the importance of a subject matter, are actually non-epistemic reasons for some activity, such as getting yourself to have detailed and accurate beliefs on that

subject matter (Way 2012). Because inquiry is often a good way to get yourself to hold detailed and accurate beliefs, this strategy posits a large number of non-epistemic reasons for inquiry which must be sharply separated from questions about rational belief.

In this section, I have developed a consequentialist extension of the reason-responsiveness approach sketched in Section 2. This view holds that all reasons for activities such as inquiry are reasons to promote value, and that rationality requires us to do what is best. In order to avoid marginalizing epistemic reasons for belief and inquiry, I adopt a rich axiology on which intellectual achievements bear significant final value, as well as a sharp level separation between the rationality of belief and inquiry.

The purpose of this account is to clarify and ground claims about rational inquiry. It is time to see whether a reason-responsiveness consequentialist approach can do the work set out for it. In Sections 4-5, I show how the view accounts for two puzzle cases in recent discussions of rational inquiry: norms of clutter-avoidance and logical non-omniscience. In Section 6, I close with a discussion of the sense in which these norms are norms of epistemic rationality.

4 Clutter avoidance

If you are like most people, you believe the following: the sky is blue. From this belief you can deduce many others. For example, either the sky is blue or today is Quine's birthday. Must you spend your days drawing out each trivial logical consequence of your beliefs?

Many philosophers have followed Gilbert Harman in thinking that you should not do this. To clutter your mind with trivialities would be an irrational waste of time. Harman captured this thought in a principle of Clutter Avoidance.

Clutter Avoidance (CA) One should not clutter one's mind with trivialities.

(Harman 1986, p. 12).

Harman gave three arguments for CA. First, *wasted effort* is expended in deducing trivialities. Second, due to our *limited retrieval capacity*, cluttering long-term memory with

irrelevant beliefs makes it more difficult to retrieve relevant beliefs from memory. Finally, our memories have *limited storage capacities* for information.

While the argument from limited storage capacities has met with mixed reactions, most philosophers have thought that there is something plausible about CA and that considerations of wasted effort and limited retrieval capacity are an integral part of the story.¹² Philosophers have disagreed about the contents of CA: what precisely does CA require? Philosophers have also disagreed about the explanation for the truth of CA. In this section, I review two prominent views about the contents of CA and its grounds, then develop an alternative consequentialist view.

Jane Friedman (2018) takes Clutter Avoidance to concern *junk beliefs*. Friedman begins with the notion of a junk subject matter. Subject matters are junk for agents whose interests would not be served by knowing about them. More precisely, subject matter x is junk for agent S in world w at time t just in case S has no interest or desire served by having a belief about some proposition contained in x . Junk propositions are members of junk subject matters. Friedman takes Clutter Avoidance to be a prohibition against forming junk beliefs.

(CA_F) Necessarily, if p is junk for S at w, t , then S ought not believe p at w, t .

Friedman accepts Harman's motivations for CA_F. She also takes on board Harman's view that CA_F is not a first-order normative claim, but rather a meta-normative constraint on acceptable norms of belief revision. By this, Friedman means that norms of belief revision which conflict with CA_F cannot be genuine norms.

Note that CA_F, like CA, is phrased as a claim about oughts rather than rational requirements. The claim that rationality is deontic explains why this is not a conceptual mistake: we can pass between the coextensive categories of what agents ought to do and what they are rationally required to do. For this reason, we can meet Friedman halfway and speak in this section about how agents ought to inquire without risk of equivocation.

¹²The trouble with the argument from storage capacities is that many memory researchers think we are in limited danger of running out of memory storage (Michaelian 2011).

While I will take on board substantial portions of Friedman's view, the view has three limitations which a consequentialist approach helps resolve. First, inquirers should not merely be concerned with their own interests but also with the interests of others.¹³ Suppose that Sharky Sue is a CEO who is genuinely unconcerned with the welfare of her workers and benefits financially from exploiting them. Should Sharky Sue inquire into her workers' welfare? According to CA_F , Sue should not. The welfare of Sue's workers is a junk subject matter for Sue, since it would serve none of Sue's interests to have beliefs about her workers' welfare. But Sue's unconcern for her workers' welfare does not absolve her from the obligation to learn about their welfare any more than it absolves Sue from the obligation to improve her workers' welfare. We have no more normative license to be selfish in inquiry than in any other activity.

Second, although wasted effort and limited retrieval capacity are compelling motivations for CA_F , we need to explain why these considerations are normatively relevant to inquiry. As Friedman notes, many 'epistemic purists' take such considerations to be irrelevant to rational inquiry. A compelling response will not merely deny the epistemic purist's contentions, but also substitute an alternative picture of rational inquiry on which these considerations are normatively relevant.¹⁴ We will see at the end of this section how consequentialism captures the normative relevance of wasted effort and limited retrieval capacities. We will also see at the end of Section 5 how consequentialism captures three further normative considerations driving theorists like Harman to adopt Clutter Avoidance.

Third, we need to clarify the objects to which CA_F applies. Although CA_F is stated as a principle governing belief states, Friedman clarifies that CA_F governs belief revision, that is the process of inquiry by which our beliefs are revised. Clutter Avoidance has

¹³This is the main point of contention between consequentialist and instrumentalist (Steglich-Petersen forthcoming) accounts.

¹⁴Friedman (2020) provides one such picture: all epistemic norms are zetetic norms, and all zetetic norms are epistemic norms. We will see in Section 6 that the Reason-Responsiveness Consequentialist View makes room for this claim, but also opens the possibility of weaker views which allow a substantive distinction to be drawn between epistemic and non-epistemic norms.

been applied to at least three separate processes of belief-revision: initial investigation, encoding of beliefs in long-term memory, and forgetting. By separately specifying the demands of Clutter Avoidance on each of these processes we can clarify its contents and motivations. In each case, I will argue that in addition to any meta-normative content, CA_F generates specific first-order normative requirements on investigation, encoding, and forgetting.¹⁵ In formulating these norms I will leave open the conditions under which a belief counts as junk, since on my own view the junkiness of a belief is not essentially tied to a believer's interests.¹⁶

Suppose you are walking down the street and you notice that p : there is a ladybug on your neighbor's rosebush. As a result, you form the junk belief that p . Have you done something wrong? Plausibly not. We form junk beliefs of this sort all the time, and it would be more effort than it is worth to prevent ourselves from forming them. But it would have been wrong for you to search the rosebush for ladybugs unless you were curious or wanted to protect the roses. This suggests that there is a clutter-avoidance norm on investigation driven by considerations of wasted effort.

Junk Non-Investigation (JNI) Agents should not expend effort investigating junk subject matters.

You have not violated JNI because you did nothing more than notice that p in passing. It would have cost you more effort to suppress the junk belief than to form it.

Now that you have formed the belief that p , should you encode that belief in long-term memory? Some philosophers think that memory encoding is involuntary and hence inapt for rational assessment.¹⁷ However, at the very least it should be conceded that

¹⁵Friedman's argument against a first-order construal of Clutter Avoidance draws on her view that Clutter Avoidance is interest-sensitive, which I deny. Friedman is also concerned that first-order Clutter Avoidance norms will conflict with existing epistemic norms, but this is not clearly the case if existing norms are taken as constraints on belief states rather than processes of inquiry, given a sharp level separation between belief and inquiry.

¹⁶For consequentialists, 'junk' is an evaluative term in the standard axiological sense of the word. Junk beliefs are beliefs without significant value or disvalue. In this usage, 'junk' is opposed to 'significant', and the junkiness or significance of a belief is a matter of degree. The turn to a graded notion of junkiness is needed to block avoid trivialization, since few beliefs are utterly without value.

¹⁷To be clear, it is not my view that only voluntary processes are apt for rational assessment. I also take

there are norm-governed voluntary processes which influence the depth and likelihood of encoding beliefs in long-term memory. For example, you can increase your probability of recalling a person's name by verbally repeating it after meeting them, and decrease the probability of successful encoding by distracting yourself. In the case of junk beliefs p , it is plausible at least that you should take cognitively inexpensive steps to ensure that p is not encoded, for example not staring at the ladybug. Otherwise p will clutter later memory searches, making it more difficult for you to make effective judgments and decisions based on information stored in memory. This motivates another first-order Clutter Avoidance norm.

Junk Non-Encoding (JNE) Agents should take steps to ensure that junk beliefs are not encoded in long-term memory.

Plausibly, you have not violated JNE if you merely notice the ladybug in passing and end up encoding the belief that p in long-term memory.

But now that p is encoded in long-term memory, all is not lost. We can and do preferentially forget junk beliefs in order to declutter memory retrieval. Here is how the trick is done (Anderson and Schooler 1991). Junk beliefs are unlikely to have been accessed frequently or recently, since they are irrelevant to most memory queries. By preferentially forgetting memories according to recall frequency and recency, we can rid ourselves of junk beliefs over time. Again, many philosophers think that memory is inapt for normative assessment, but this opinion is increasingly controversial: we regularly criticize agents for failures to remember (Carr 2015; Smith 2005) and a growing number of philosophers and psychologists develop rationality requirements on memory (Schooler and Hertwig 2005; Bernecker 2018).

I want to side-step the question of which steps towards forgetting are sufficiently volitional to be apt for rational assessment. Most philosophers accept that some forgetting-related steps can be the subject of rational requirements. For example, we can be rationally

a relatively broad view of the class of memory processes that should be understood as voluntary.

required not to intentionally mentally replay an unfortunate episode that we would like to forget, because replaying the episode decreases the likelihood that we will forget it. Hence without taking a stand on the extent of rational requirements on forgetting, we can see Clutter Avoidance as imposing the following norm:

Junk Forgetting (JF) Agents should take steps to ensure that junk beliefs are forgotten.

If you are like most people, you will soon forget all about the ladybug on the rosebush. That is a good sign that you have complied with JF.

Summing up, we need to develop an account of Clutter Avoidance which is unselfish, explains why standard motivations for Clutter Avoidance are good ones, and separately accounts for the first-order normative requirements imposed by Clutter Avoidance. A good start at this project is made by Kourken Michaelian (2011). Michaelian argues that the function of memory is to provide a manageable quantity of currently relevant information in a timely manner. These considerations could be taken to directly support JF and JNE, although JNI would need to be separately motivated by appeal to wasted effort. But this is not what Michaelian does. Michaelian argues that Clutter Avoidance can be supported by a purely intellectual norm which ranks belief states by the ratio of junk to non-junk beliefs contained in them. Roughly, the norm is:

Ratio (RT) Let B, B' be belief states which differ only in that B' contains some junk beliefs that B does not contain. Then B' is epistemically inferior to B .

Michaelian motivates RT by arguing that we care intrinsically about the ratio of junk to non-junk beliefs contained in a belief state. I must confess that I find it opaque why agents should care about such a thing. It is perfectly intelligible why agents would be averse to spending effort moving from B to B' , or why they would be worried that B' clutters later memory retrieval. It is also intelligible why agents would think that B' is not much epistemically better than B , since the junk beliefs in B' are not very important. But why should the ratio of junk to non-junk beliefs contained in an agent's belief state count

intrinsically against it? RT is an ad-hoc principle invented to deliver Clutter Avoidance norms that are best explained by the very arguments that Harman took to motivate them.

A consequentialist approach to Clutter Avoidance fares better. The Reason-Responsiveness Consequentialist View explains the normative importance of limited retrieval capacities. As Michaelian correctly notes, in memory retrieval agents aim to retrieve manageable quantities of relevant information in a timely manner. We do this in order to make good decisions and form correct judgments on the basis of information retrieved, because on a rich axiology it is better to make good decisions and judgments than bad ones. The need to declutter retrieval motivates JNE and JF.

Considerations of wasted effort are also normatively relevant to our initial investigations. It is wrong to waste effort investigating junk subject matters, even if the resulting beliefs will not be encoded in long-term memory. Wasting effort is wrong because that effort could be better spent on other pursuits. Limited retrieval capacities are also relevant insofar as junk investigations may, against our wishes, lead junk beliefs to be encoded in long-term memory, cluttering retrieval. Together, wasted effort and limited retrieval capacities provide good support for JN1.

In this way, the Reason-Responsiveness Consequentialist View explains the normative relevance of Harman and Michaelian's original motivations for Clutter Avoidance norms. Consequentialism grounds the correct first-order normative consequences of Clutter Avoidance, and does so without introducing unmotivated epistemic principles. Further, consequentialism is unselfish. Sharky Sue ought to investigate and take steps to remember facts about her workers' welfare because knowing these facts puts Sue in a position to improve the lives of her workers.

We have seen that the Reason-Responsiveness Consequentialist View explains the nature and grounds of Clutter Avoidance norms, capturing the most plausible claims made by existing views while avoiding their defects. In the next section, we will see how the Reason-Responsiveness Consequentialist View solves a second puzzle: accounting for norms of logical non-omniscience.

5 Logical omniscience

How does logic constrain rational inquiry? On some views, rational inquirers are subject to norms of logical omniscience such as the following:¹⁸

(Deductive Closure) If a rational agent believes that p_1, \dots, p_n , and p_1, \dots, p_n entail q , then she forms the belief that q or abandons her belief in one of the p_i 's.

(Consistency) If a rational agent believes an inconsistent collection of claims p_1, \dots, p_n , then she abandons her belief in one of the p_i 's.

Authors studying rational inquiry are nearly unanimous in their opposition to norms of logical omniscience (Harman 1986; Hintikka 2001).

At least three reasons are given for rejecting norms such as Deductive Closure and Consistency. First, full compliance is *beyond our abilities*. We cannot, for example, deduce all logical consequences of Peano Arithmetic. Second, it is often *too demanding* to comply with omniscience norms even when we could do so. To infer from p to $p \vee q$ is within our abilities, but often a waste of time. In fact, Gilbert Harman took such inferences as his original motivation for Clutter Avoidance. Finally, as Harman (1986) notes, the laws of logic are *principles of implication*. They tell us what follows from what. Principles of implication are not *principles of revision*, telling us how to revise our beliefs. We can propose principles of revision which are motivated by the laws of logic, but the status of these logical laws as principles of implication is no reason to accept them as principles of belief revision. Indeed, it would be surprising if there were a straightforward correspondence between the very different questions of what follows from what, and which processes rational agents should use to revise their beliefs.

¹⁸MacFarlane (manuscript) and Steinberger (2019) term these 'bridge principles'. While a full discussion of their views is beyond the scope of this paper, the move towards attitudinal restriction in Steinberger (2019) will be covered later in this section. Some other important issues, such as the distinction between narrow- and wide-scope readings, go beyond the focus of this paper.

It might be thought that weakened forms of Deductive Closure and Consistency survive these criticisms. For example, we could restrict both principles by requiring that the entailments in question be foreseeable or beneath a certain logical complexity. But such weakenings are at once too demanding and not demanding enough. These weakenings are too demanding because simple and foreseeable inferences such as the inference from p to $p \vee q$ may do nothing but waste effort and contribute to mental clutter. And they are not demanding enough because they make no room for stakes. For example, humans make a variety of logical errors in evaluating important financial decisions, and these errors are enormously costly (Benartzi and Thaler 2007; Johnson et al. 1993). We could reduce the frequency of such errors by employing more demanding forms of reasoning which draw less foreseeable and more complex inferences and detect or avoid additional inconsistencies. But nothing in the weakened forms of Deductive Closure or Consistency explains why we ought to do so.¹⁹

Some authors respond by saying too little about the logical requirements on rational inquiry. For example, Christopher Cherniak weakens Consistency to say that:

If an agent has a particular belief-desire set, then if some (but not necessarily all) inconsistencies arose in his belief set, he would eliminate them. (Cherniak 1981, p. 172).

This is not false, but it does not tell us what we would like to know. We want to know which inconsistencies a rational inquirer will detect and eliminate. Similarly, Harman weakens Consistency to hold:

One has a reason to avoid believing things one recognizes to be inconsistent.
(Harman 1986, p. 18).²⁰

¹⁹It might be objected that the threshold for complexity or foreseeability is context-sensitive. But now we need an account of how and why these thresholds vary, and here consequentialism seems a likely candidate.

²⁰For a similar principle see Steinberger (2019). Note that Steinberger takes this principle to be a *directive*, offering first-personal guidance, rather than an *evaluative* third-person standard. If that is right, then we need not disagree since consequentialism is an evaluative claim. However, Steinberger later claims that the demandingness of inquiry is not evaluatively relevant. Here we do disagree.

Again, this is too weak. We want to know which inconsistencies a rational inquirer should recognize, not merely that she has a reason to rid herself of inconsistencies once they are recognized.

At this point, we might despair of saying anything substantive and true about the logical requirements on rational inquiry. But that would be a mistake. To see the way forward, consider the game of chess. Most chess players evaluate moves at least partly by considering candidate continuations, sequences of future moves that could result from a given position. Logically omniscient agents would not need to generate and examine novel continuations, since the existence of each continuation follows by first-order logic from the agent's knowledge of the present position and the rules of chess. But for non-omniscient agents, generating new continuations is a way of increasing the deductive closure of their belief state. Here we can ask the natural normative question: when should agents consider additional continuations, and which move sequences should agents search for continuations of?

These questions can be most precisely posed for agents whose computational abilities and memory are fully specified. A great deal of work has been done on these questions. Some answers on offer are explicitly consequentialist. For example, agents might engage in consequentialist meta-reasoning, assigning to each position that they are aware of an expected value of searching for continuations beginning at that position and taking the searches with highest expected value (Lieder and Griffiths 2017; Russell and Wefald 1991). The expected value of search tracks the expected improvement to the player's position, and perhaps also intellectual achievements such as gained understanding, less the expected cost of deliberation time. Most authors hold that explicitly calculating expectations is more costly than it is worth. These authors instead propose algorithms for pruning and extending trees of possible continuations.²¹ Although these algorithms do not explicitly calculate expectations, most scholars have understood them to have an

²¹An accessible and effective example is alpha-beta pruning in minimax decisionmaking (Knuth and Moore 1953). Some recent algorithms function quite differently (Silver et al. 2018).

explicitly consequentialist justification. If agents ought to evaluate some continuations rather than others or halt search rather than continue, that is because doing so will have the best impact on the player's odds of winning the game or their understanding of chess play.

I take this discussion to reveal two things. First, there is often a great deal that can be said about logical requirements on rational inquiry. In the special context of evaluating possible continuations during chess play, there is a well-established body of research offering specific and detailed characterizations of the degree of deductive closure that agents should achieve given constraints imposed by their cognitive capacities and deliberation time. Second, at least in the special case at hand, norms of inquiry are explicitly consequentialist: agents should inquire in the ways that are expectedly best, which in this case amounts to doing whatever makes them most likely to win the game.²² Features such as foreseeability and complexity matter only insofar as they bear on the possibility and cost of examining some continuations rather than others. There are many foreseeable combinations, and even single moves which agents may rationally fail to consider.

Consequentialists think that logic plays the same normative role in chess as it does in all other inquiries. We can ask various questions about rational principles of revision. When should inferences be made using logical principles such as modus ponens? When may agents use inference patterns which sometimes produce inconsistent beliefs, and how much potential inconsistency should they accept? Each of these questions is to be answered on consequentialist grounds by citing the expected improvements to an agent's belief state, the importance of those improvements and the costs of making them. If they wish, consequentialists can hold that coherence, like accuracy, has intrinsic value, but the intrinsic value of coherence is not always the only, or even the most relevant normative factor (Thorstad forthcoming a).

²²In taking expectations, I side with many epistemologists against objectivists who think that agents should do what will actually have the best consequences. I will not have the time to argue for this view here. Objectivists are invited to repeat this discussion without taking expectations.

To see how these questions generalize beyond the artificial context of chess play, consider the problem of strategy selection during inquiry (Marewski and Schooler 2011; Lieder and Griffiths 2017). Throughout our lives we must choose between frugal, heuristic forms of inquiry and more demanding nonheuristic forms of inquiry (Thorstad 2022b, forthcoming a,c). In many cases, heuristics are less costly than nonheuristic inquiries, but are slightly less accurate and carry a small risk of producing incoherent judgments. As bounded agents, humans must sometimes inquire heuristically, but we needn't always do so.²³ When are the costs of heuristic inquiry, in particular the risk of incoherence, worth paying?

Again, it will help to consider an example. Suppose I ask you to judge whether a distant tree is an oak tree. You might make this judgment by considering how representative the distant tree is of a typical oak tree. Does it have acorns? Are its height and shape typical of an oak tree? Here you have employed the *representativeness heuristic*, determining whether an object, the distant tree, is likely to belong to a category, oak trees, by asking how representative the object is of the category. The representativeness heuristic is often a quick and reliable way of making inferences, but it differs from nonheuristic forms of inference in a number of ways. For example, the representativeness heuristic neglects base rates. Nowhere in your inference did you consider the prevalence of oak trees in your surroundings. In some contexts, this leads representativeness-based inferences to be wildly inaccurate, but in many contexts the drop in accuracy may be mild and well worth paying.

However, in addition to their reduced accuracy, representativeness-based inferences are occasionally incoherent. The problem is easiest to see in the context of probabilistic inference. Daniel Kahneman and Amos Tversky gave participants the following vignette:

Linda is 31 years old, single, outspoken, and very bright. She majored in

²³This claim is less controversial than it may appear. For example, Tversky and Kahneman write: "It is not surprising that useful heuristics such as representativeness and availability are retained, even though they occasionally lead to errors in prediction or estimation" (Tversky and Kahneman 1974, p. 1130), and the introduction to a recent anthology in the heuristics and biases tradition clarifies: "heuristics themselves are sensible estimation procedures that are by no measure 'irrational' " (Gilovich and Griffin 2002, p.3).

philosophy. As a student, she was deeply concerned with the issue of discrimination and social justice, and also participated in antinuclear demonstrations. (Tversky and Kahneman 1983).

Kahneman and Tversky asked one group of participants to estimate the likelihood that Linda is a bank-teller, and another group to estimate the likelihood that Linda is a feminist bank-teller. Participants in aggregate made the incoherent judgment that Linda is more likely to be a feminist bank-teller than a bank-teller. The explanation for this result is that agents made their judgments using the representativeness heuristic. Linda is more representative of a feminist bank-teller than a bank-teller, so participants judged that Linda was more likely to be a feminist bank-teller than a bank-teller.

As our discussion of representativeness illustrates, heuristics are often less costly than nonheuristic inquiries, but are slightly less accurate and carry a small risk of producing incoherent judgments. When are these costs worth paying? Consequentialists think that potential incoherence should be balanced against cognitive frugality in exactly the same way as frugality is balanced against potential inaccuracy. When the potential for incoherence or inaccuracy is sizable, or when these consequences are especially important to avoid, agents should not inquire heuristically. When these costs are bearable, agents should inquire heuristically in order to keep the costs of inquiry manageable.

The consequentialist account has a number of virtues. Chief among them, it lets us see how Linda is not like an oak tree.²⁴ There may be nothing wrong with using representativeness to determine whether a distant tree is an oak, but it seems seriously wrong to judge Linda's profession by representativeness alone. Why is this so? In the context of social inference, the representativeness heuristic is a simple form of stereotyping. Stereotyping is not only wrong because it produces inaccurate judgments: some stereotypes are reliable. Nor is the wrongness of stereotyping exhausted by the fact that it occasionally produces incoherent beliefs. One thing that is wrong with stereotyping is that it produces

²⁴The discussion in this section is, by necessity, somewhat compressed. For a fuller discussion of the prospects for a consequentialist account of stereotyping see Rinard (2019), Thomsen (2011) and the exchange between Risse and Zeckhauser (Risse and Zeckhauser 2004; Risse 2007) and Lever (Lever 2005, 2007).

prejudiced beliefs about Linda. Social inferences made on the basis of stereotypes cause systematic and pervasive harms to the lives of stereotyped groups, and these harms bear directly on the rationality of stereotyping as a principle of belief revision.

Epistemologists increasingly accept the need for norms of inquiry to capture the wrongness of stereotyping, as well as to make room for the permissibility of seemingly equivalent inferences, as in our example about oak trees (Basu 2019; Bolinger 2020). The Reason-Responsiveness Consequentialist View is in a position to do both, since for the consequentialist all harms caused by stereotyping bear directly on its normative status as a principle of belief revision. Here consequentialism takes a balanced view of the normative importance of incoherence and points us instead to what matters most: not incoherence, nor always inaccuracy, but rather the harms imposed on agents by stereotyping.

Summing up, we need to provide an account of the logical requirements on rational inquiry. Existing accounts will not do, but the Reason-Responsiveness Consequentialist View offers good hope for a detailed and specific account. Consequentialism explains clearly how agents ought to close their beliefs under logical deduction in a special case drawn from chess play, and extends the same account to cover all types of inquiry. The consequentialist account sheds light on a pressing question: when should agents employ heuristic rather than nonheuristic forms of inquiry? Here consequentialism suggests that potential incoherence should be balanced against cognitive frugality in exactly the same way as frugality is balanced against inaccuracy. This account helps us to see why many social inferences made on the basis of stereotypes are irrational.

Finally, the Reason-Responsiveness Consequentialist View captures the primary grounds on which logical omniscience requirements have been challenged: demandingness, cognitive abilities, and the distinction between principles of implication and principles of belief revision. For the consequentialist, principles of belief revision are to be justified by the costs and benefits of employing them, and that is an entirely different question from which principles of implication hold between propositions. The demandingness of principles of belief revision bears directly on their costs, and hence on their

rationality. And rational requirements cannot be beyond our abilities, since rationality is deontic and in deontic matters, ought implies can. Here, as before, consequentialism provides a well-motivated, explanatorily adequate account of proposed norms of inquiry. That account is relatively non-revisionary, both in its normative recommendations and in its ability to capture existing motivations for proposed norms. These results support a Reason-Responsiveness Consequentialist View of the nature and grounds of norms of inquiry.

6 Inquiry and the epistemic

Some readers may think that my account succeeds by changing the subject. The Reason-Responsiveness Consequentialist View is an account of rationality simpliciter. But, the objection continues, the project was not to characterize rationality simpliciter, but rather to characterize a special type of epistemic rationality which applies to inquiry. My view carries out the wrong project, because it is not an account of epistemic rationality. Call this the *epistemic objection*.

In this section, I develop two lines of response to the epistemic objection: a soft response (Section 6.1) and a strong response (Section 6.2). Although my own sympathies lie with the strong response, either response would be sufficient to motivate the truth and importance of the Reason-Responsiveness Consequentialist View.

6.1 The soft response

In this paper, I have given accounts of a number of normative notions: rational inquiry, rationally permissible inquiry, rationally required inquiry, and how agents should or ought to inquire. Normative concepts such as rationality and ought are, by any reckoning, among the most central and important normative concepts in philosophy. These concepts also play an important role in our everyday practices of normative evaluation.

Some readers may think that there are additional concepts, such as epistemic ratio-

nality, worth theorizing about. For that matter, it may also be important to think about fittingness (McHugh and Way 2016), blameworthiness (Brown 2020), or the aim of inquiry (Falbo forthcoming; Kelp 2021b). But the existence of a special type of epistemic rationality would not detract from the importance of getting clear on rationality simpliciter, any more than the existence of fittingness and blame detracts from the importance of theorizing about epistemic rationality.

Even in the case of belief, there have long been consequentialists, instrumentalists, and others who make central appeal to traditionally non-epistemic factors in their accounts of rationality (McCormick 2015; Meylan 2021; Rinard 2019; Steglich-Petersen and Skipper 2019). The standard reaction to these accounts has not been that they are uninteresting. Indeed, it is generally thought to be quite a serious matter to settle the truth of such accounts. The standard reaction to these accounts has rather been that they are untrue: non-epistemic reasons could not be genuine reasons for belief because they are not motivating reasons (Kelly 2002, 2003), play the wrong role in deliberation (Shah 2003, 2006), or are the wrong kind of reasons for belief, and hence really reasons for some activity such as getting ourselves to believe (Way 2012). Nothing in my account forces us to abandon this reaction. By positing a strict level separation between questions about rational belief and inquiry, my account allows us to admit non-epistemic reasons for inquiry without forcing us to posit non-epistemic reasons for belief (Thorstad 2021).

However, we will see shortly that many of the same arguments that caused theorists to deny the existence of non-epistemic reasons for belief cannot be used to deny the existence of non-epistemic reasons for inquiry. If that is right, then we may well want to recognize a central role for traditionally non-epistemic reasons in rationalizing inquiry. In that case, it will be both correct and important to get a good handle on the normative impact of non-epistemic reasons for inquiry, because it is important to arrive at a correct account of rationality.

This, then, is the *soft response* to the epistemic objection: the existence of a distinctively epistemic type of rationality governing inquiry need not detract from the interest, cogency

or importance of theorizing about rational inquiry, nor tell against the relevance of traditionally non-epistemic factors to rational inquiry. Readers who are content with the soft response are welcome to rest here. But I want to push my luck a bit and test the prospects for a stronger response to the epistemic objection.

6.2 The strong response

The epistemic objection makes at least four assumptions: that there is a distinctive type of epistemic rationality which applies to inquiry; that traditionally non-epistemic factors such as consequences do not bear on the epistemic rationality of inquiry; that central normative requirements on inquiry such as clutter avoidance can be captured as requirements of epistemic rationality; and that the most important project for epistemologists studying inquiry is describe the nature of epistemically rational inquiry. All of these assumptions have been questioned by recent discussions.²⁵

Begin with the existence of a distinctive type of epistemic rationality governing inquiry. Until very recently, many if not most epistemologists held that no distinctive type of epistemic rationality applies to inquiry (Dougherty 2014; Feldman 2002; Hedden 2015), and the highest-profile exception (Hall and Johnson 1998) was met with heavy criticism (Feldman 2002). More recently, at least two epistemologists have argued against the existence of a distinctive type of epistemic rationality governing inquiry (Thorstad 2022a; Steglich-Petersen forthcoming). While there are some extended arguments for the existence of a distinctive type of epistemic rationality governing inquiry (Friedman 2020; Woodard and Flores forthcoming), there are fewer such arguments than many readers may expect, and it is by no means a foregone conclusion that these arguments will be successful.²⁶

Turn next to the irrelevance of consequences to rational inquiry. Many theorists, including central figures in the epistemology of inquiry, have explicitly held that consequences

²⁵For a full treatment of these important issues, see Thorstad (2022a, forthcoming b).

²⁶It is also not clear whether these arguments generate a distinct type of epistemic *rationality* governing inquiry, as opposed to some other normative notion (Thorstad 2023).

and other traditionally non-epistemic factors bear on the rationality of inquiry (Friedman 2020; Harman 2004; Lord 2020; Thorstad 2022a). At the most extreme, Jane Friedman (2020) holds that all norms of inquiry are epistemic, specifically mentioning instrumental norms as among the norms of inquiry that might count as epistemic. Defenders of the epistemic objection are welcome to argue that consequences should not bear on rational inquiry, at least not in the way that the consequentialist suggests, but this cannot be simply assumed.

Now let us ask whether central requirements such as clutter avoidance can be captured using within a purely epistemic approach to rationality. While some theorists have tried to do this (Michaelian 2011), others have taken a different tack. For example, we saw that Jane Friedman (2018) gives an instrumentalist account of clutter avoidance. It will be hard to find principled grounds to classify consequentialism as a non-epistemic theory without also ruling out instrumentalism. More generally, traditional epistemic accounts of clutter avoidance rely on the view that some propositions are intrinsically more epistemically interesting or valuable than others. Although it may be possible to develop a satisfactory account of epistemic interestingness or value, this project has proven more challenging than might be expected (Treanor 2014, 2018). We need to be open to the idea that an account which appeals to non-epistemic notions of value will generate accounts of clutter avoidance and other normative requirements that outperform traditionally epistemic accounts on theoretical virtues such as parsimony, explanatoriness, and theoretical motivation.

Finally, consider the view that the most important project for epistemologists is to deliver an account of epistemically rational inquiry. Although this claim may seem obvious or even definitional, several authors have denied it. I have argued (Thorstad 2022a), following Thomas Kelly (2002; 2003) that a central project for epistemology is to characterize the notion of theoretical rationality. We argue that theoretical rationality encompasses both belief and inquiry. Even if the rationality of belief is a purely epistemic matter, we argue, rational inquiry may answer to a lively mixture of epistemic and non-

epistemic reasons. If that is right, then insofar as epistemologists are concerned to capture the notion of theoretical rationality, one aim of epistemology should be to characterize the rationality of inquiry in a way that captures the relevance of epistemic and non-epistemic reasons alike.

So far, we have seen that the epistemic objection assumes at least four things: that there is a distinctive type of epistemic rationality which applies to inquiry; that traditionally non-epistemic factors such as consequences do not bear on the epistemic rationality of inquiry; that central normative requirements on inquiry such as clutter avoidance can be captured using purely a purely epistemic notion of rationality; and that the most important project for epistemologists studying inquiry is to capture the notion of epistemically rational inquiry. The *strong response* to the epistemic objection suggests that some or all of these assumptions are false.²⁷ Above, we saw how each of the assumptions made by the epistemic objection may be challenged. If these assumptions are false, then the epistemic objection may not get off the ground.²⁸

From a certain traditional standpoint, denying these assumptions can seem heretical, unmotivated, and perverse. But in fact, I think that even many theorists who are sympathetic to the epistemic objection in the case of rational belief should be prepared to question its underlying assumptions when we turn from belief to inquiry. The reason for this is that many of the same considerations which led us to posit and privilege a notion of epistemic rationality governing belief not only fail to motivate the existence or importance

²⁷How does the strong response relate to the soft response? The strong response gives reasons to doubt all four assumptions underlying the epistemic objection. By contrast, the soft response focuses only on a single assumption, that the most important project for epistemologists studying inquiry is describe the nature of epistemically rational inquiry. Unlike the strong response, the soft response does not deny this assumption, but notes that it is compatible with taking the study of rational inquiry to be an interesting and important project for epistemologists. Thanks to a referee for pressing me to clarify this matter.

²⁸Of course, to say that the assumptions behind the epistemic objection are matters of active debate is not to say that the assumptions are false. As always, readers are encouraged to make up their own minds through a detailed examination of existing arguments, since constraints of space are not sufficient to permit a full survey of all arguments in the present paper, and since many of these issues are well-treated by existing discussions. If the existing arguments tell strongly in favor of all assumptions behind the epistemic objection, then the right thing to conclude will be that the epistemic objection is correct and that this paper is rather less interesting than it may have appeared. However, if the existing arguments do not tell strongly in favor of the epistemic objection, that may be good reason to take the reason-responsiveness consequentialist view seriously. Thanks to a referee for pressing me to clarify the status of this discussion.

of a distinctive type of epistemic rationality governing inquiry, but in fact may well tell against it.

For example, consider a traditional argument due to Thomas Kelly (2002; 2003). Kelly argues that there could not, even in principle, be non-epistemic reasons for belief, because reasons for belief must be potential motivating reasons for belief, and only evidence can motivate belief. On this basis, Kelly argues that defenders of non-epistemic reasons for belief commit the *consequentialist mistake* of assuming that since consequences bear on the rationality of actions, they also bear on the rationality of inquiry. This is a mistake because consequences can be motivating reasons for action, but not for belief.

However, following Kelly (2002; 2003), I have emphasized (Thorstad 2022a) that the same is not true of inquiry. Non-epistemic reasons can be, and frequently are motivating reasons for inquiry. For example, you might check the train schedule so that you will not be late to work. Because of this, I have argued (Thorstad 2022a) that denying or marginalizing non-epistemic reasons for inquiry would commit the *reverse-consequentialist mistake* of assuming that since consequences do not bear on the rationality of belief, they also do not bear on the rationality of inquiries which produce belief. This is a mistake because although consequences cannot be motivating reasons for belief, they can be motivating reasons for actions such as inquiry.

More generally, I have argued (2022a; forthcoming b) that many of the same reasons which led us to posit and privilege a distinctive type of epistemic rationality governing belief should, if anything, make us more reluctant to posit and privilege a distinctive type of epistemic rationality governing inquiry. If this is right, then there need not be anything shocking, heretical or counterintuitive about the Reason-Responsiveness Consequentialist View, which is an account of rational inquiry rather than an account of rational belief. The view may instead correctly respond to the failure of traditional motivations for a positing a distinctive type of epistemic rationality to generalize from belief to inquiry.

In this section, we considered the epistemic objection to the Reason-Responsiveness Consequentialist View. The epistemic objection holds that the view answers the wrong

question because it delivers an account of rationality simpliciter rather than epistemic rationality. We explored two responses to the epistemic objection. The soft response held that even if there exists a distinctive type of epistemic rationality which applies to inquiry, it remains very important to characterize the general notion of rational inquiry. The strong response showed how a number of assumptions made by the epistemic response may be resisted. Either of these responses would be enough to ground the truth and importance of the Reason-Responsiveness Consequentialist View.

7 Conclusion

My aim in this paper was to develop a general account of rational inquiry and to apply the account to clarify and ground two types of normative requirements governing inquiry. I began with a reason-responsiveness account of rationality (Section 2) and extended that account with a consequentialist account of normative reasons for inquiry (Section 3) to generate a reason-responsiveness, consequentialist view of rational inquiry, the Reason-Responsiveness Consequentialist View. I applied this account to clarify and ground norms of clutter avoidance (Section 4) and requirements of logical non-omniscience (Section 5). This discussion gave rise to the epistemic objection that the Reason-Responsiveness Consequentialist View answers the wrong question, because it develops an account of rationality simpliciter rather than epistemic rationality. I developed two lines of response to the epistemic objection in Section 6.

This discussion is preliminary in several ways. First, while I have considered two case studies of specific norms governing inquiry (Sections 3-4), there are many other normative questions about inquiry which I have not considered. For example, we might ask when and how inquirers should gather evidence (Hall and Johnson 1998; Woodard and Flores forthcoming), double-check their beliefs (Friedman 2019; Woodard forthcoming), or allocate attention during inquiry (Siegel 2017). Traditionally, consequentialist accounts have earned their keep by delivering plausible, well-motivated and precise accounts of

a wide range of normative requirements. By exploring what the Reason-Responsiveness Consequentialist View has to say about these and other norms governing inquiry, we will be in a better position to understand the contents and implications of the view. If those implications are plausible, that will be some evidence in favor of the Reason-Responsiveness Consequentialist View, and if they are implausible, then we may have reason to modify the view.

Second, we saw in Section 6 that there is room for further discussion of the nature and existence of a distinctively epistemic type of rationality governing inquiry, as well as the relationship between rational inquiry and epistemically rational inquiry. Although these questions have been the object of some recent discussion (Friedman 2020; Steglich-Petersen forthcoming; Thorstad 2022a), I hope that future discussions will help us to reach further clarity on the range of viable views, and the central considerations for and against these views.

Finally, the Reason-Responsiveness Consequentialist View is not the only possible account of rational inquiry. With the possible exception of Kelp (2021a,b), recent discussions of rational inquiry have not produced any competing accounts of rational inquiry.²⁹ It would be helpful to have more accounts on the table in order to understand the space of possible views, and to compare the performance of competing accounts across a range of specific normative requirements such as clutter avoidance, evidence gathering, and norms of logical non-omniscience. It is my hope that the discussion in this paper will provide useful input into the construction and assessment of alternative accounts.

References

Ahlstrom-Vij, Kristoffer and Dunn, Jeffrey (eds.). 2018. *Epistemic consequentialism*. Oxford University Press.

²⁹I say that this is a possible exception because I am not sure whether Kelp intends his analysis be an account of rationality, or of some other normative notion.

- Anderson, John and Schooler, Lael. 1991. "Reflections of the environment in memory." *Psychological Science* 2:396–408.
- Arpaly, Nomy and Schroeder, Timothy. 2015. "A causal theory of acting for reasons." *American Philosophical Quarterly* 52:103–14.
- Basu, Rima. 2019. "Radical moral encroachment: The moral stakes of racist beliefs." *Philosophical Issues* 29:9–23.
- Benartzi, Shlomo and Thaler, Richard. 2007. "Heuristics and biases in retirement savings behavior." *Journal of Economic Perspectives* 21:81–104.
- Berker, Selim. 2013a. "Epistemic teleology and the separateness of propositions." *Philosophical Review* 122:337–93.
- . 2013b. "The rejection of epistemic consequentialism." *Philosophical Issues* 23:363–387.
- Bernecker, Sven. 2018. "On the blameworthiness of forgetting." In Kourken Michaelian, Dorothea Debus, and Denis Perrin (eds.), *New directions in the philosophy of memory*, 241–58. Routledge.
- Bolinger, Renée. 2020. "The rational impermissibility of accepting (some) racial generalizations." *Synthese* 197:2415–31.
- Brown, Jessica. 2020. "What is epistemic blame?" *Noûs* 54:389–407.
- Carr, Jennifer. 2015. "Subjective ought." *Ergo* 2:678–710.
- Cherniak, Christopher. 1981. "Minimal rationality." *Mind* 90:161–183.
- Dougherty, Trent. 2014. "The ethics of belief is ethics (period): Reassigning responsibility." In Jonathan Matheson and Rico Vitz (eds.), *The ethics of belief*, 146–68. Oxford University Press.
- Driver, Julia. 2001. *Uneasy Virtue*. Cambridge University Press.

Falbo, Arianna. forthcoming. "Inquiring minds want to improve." *Australasian Journal of Philosophy* .

Feldman, Richard. 2002. "Epistemological duties." In Paul Moser (ed.), *The Oxford handbook of epistemology*, 362–83. Oxford University Press.

Friedman, Jane. 2017. "Inquiry and belief." *Noûs* 53:296–315.

—. 2018. "Junk beliefs and interest-driven epistemology." *Philosophy and Phenomenological Research* 97:568–83.

—. 2019. "Checking again." *Philosophical Issues* 29:84–96.

—. 2020. "The epistemic and the zetetic." *Philosophical Review* 129:501–36.

—. forthcoming. "The aim of inquiry?" *Philosophy and Phenomenological Research* .

Gibbard, Allan. 2003. *Thinking how to live*. Harvard University Press.

Gilovich, Thomas and Griffin, Dale. 2002. "Heuristics and biases: Then and now." In Thomas Gilovich, Dale Griffin, and Daniel Kahneman (eds.), *Heuristics and biases: The psychology of intuitive judgment*, 1–18. Cambridge University Press.

Greco, Daniel and Hedden, Brian. 2016. "Uniqueness and metaepistemology." *Journal of Philosophy* 113:365–95.

Hall, Richard and Johnson, Charles. 1998. "The epistemic duty to seek more evidence." *American Philosophical Quarterly* 35:129–39.

Harman, Gilbert. 1986. *Change in view*. MIT Press.

—. 2004. "Practical aspects of theoretical reasoning." In Alfred R. Mele and Piers Rawling (eds.), *The Oxford handbook of rationality*, 45–56. Oxford University Press.

Hedden, Brian. 2015. "Time-slice rationality." *Mind* 124:449–91.

Hintikka, Jaakko. 2001. "Is logic the key to all good reasoning?" *Argumentation* 15:35–57.

Horowitz, Sophie and Dogramaci, Sinan. 2016. "An argument for uniqueness about evidential support." *Philosophical Issues* 26:130–47.

Johnson, Eric, Hershey, John, Meszaros, Jacqueline, and Kunreuther, Howard. 1993. "Framing, probability distortions, and insurance decisions." *Journal of Risk and Uncertainty* 7:35–51.

Kagan, Shelly. 2000. "Evaluative focal points." In Brad Hooker, Elinor Mason, and Dale E. Miller (eds.), *Morality, rules and consequences*, 134–55. Edinburgh University Press.

Kelly, Thomas. 2002. "The rationality of belief and some other propositional attitudes." *Philosophical Studies* 110:163–96.

—. 2003. "Epistemic rationality as instrumental rationality: A critique." *Philosophy and Phenomenological Research* 66:612–40.

Kelp, Christoph. 2021a. *Inquiry, knowledge, and understanding*. Oxford University Press.

—. 2021b. "Theory of inquiry." *Philosophy and Phenomenological Research* 103:359–84.

Kiesewetter, Benjamin. 2017. *The normativity of rationality*. Oxford University Press.

Knuth, Donald and Moore, Ronald. 1953. "An analysis of alpha-beta pruning." *Artificial Intelligence* 6:293–326.

Kolodny, Niko. 2005. "Why be rational?" *Mind* 114:509–63.

Lever, Annabelle. 2005. "Why racial profiling is hard to justify: A response to Risse and Zeckhauser." *Philosophy and Public Affairs* 33:94–110.

—. 2007. "What's wrong with racial profiling? Another look at the problem." *Criminal Justice Ethics* 26:20–8.

Levinstein, Ben. 2017. "Permissive rationality and sensitivity." *Philosophy and Phenomenological Research* 94:342–70.

- Lieder, Falk and Griffiths, Thomas. 2017. "Strategy selection as rational metareasoning." *Psychological Review* 124:762–94.
- Lord, Errol. 2017. "What you're rationally required to do and what you ought to do (are the same thing!)." *Mind* 126:1109–54.
- . 2018. *The importance of being rational*. Oxford University Press.
- . 2020. "Suspension of judgment, rationality's competition, and the reach of the epistemic." In Sebastian Schmidt and Gerhard Ernst (eds.), *The ethics of belief and beyond*, 126–45. Routledge.
- MacFarlane, John. manuscript. "Is logic normative for thought?" Manuscript.
- Marewski, Julian and Schooler, Lael. 2011. "Cognitive niches: An ecological model of strategy selection." *Psychological Review* 118:393–437.
- McCormick, Miriam. 2015. *Believing against the evidence: Agency and the ethics of belief*. Routledge.
- McGrath, Matthew. 2021. "Being neutral: Suspension of judgement, agnosticism and inquiry." *Noûs* 55:463–84.
- McHugh, Conor and Way, Jonathan. 2016. "Fittingness first." *Ethics* 126:575–606.
- Meylan, Anne. 2021. "Doxastic divergence and the problem of comparability. Pragmatism defended further." *Philosophy and Phenomenological Research* 103:199–216.
- Michaelian, Kourken. 2011. "The epistemology of forgetting." *Erkenntnis* 74:399–424.
- Parfit, Derek. 1984. *Reasons and persons*. Oxford University Press.
- Pettit, Philip and Smith, Michael. 2000. "Global consequentialism." In Brad Hooker, Elinor Mason, and Dale E. Miller (eds.), *Morality, rules and consequences*, 121–133. Edinburgh University Press.

- Rinard, Susanna. 2019. "Equal treatment for belief." *Philosophical Studies* 176:1923–50.
- Risse, Mathias. 2007. "Racial profiling: A reply to two critics." *Criminal Justice Ethics* 26:4–19.
- Risse, Mathias and Zeckhauser, Richard. 2004. "Racial profiling." *Philosophy and Public Affairs* 32:131–70.
- Russell, Stewart and Wefald, Eric. 1991. "Principles of metareasoning." *Artificial Intelligence* 49:361–95.
- Sapir, Yasha and van Elswyk, Peter. 2021. "Hedging and the ignorance norm on inquiry." *Synthese* 199:5837–59.
- Schooler, Lael J. and Hertwig, Ralph. 2005. "How forgetting aids heuristic inference." *Psychological Review* 112:610–28.
- Shah, Nishi. 2003. "How truth governs belief." *Philosophical Review* 112:447–82.
- . 2006. "A new argument for evidentialism." *Philosophical Quarterly* 56:481–98.
- Siegel, Susanna. 2017. *The rationality of perception*. Oxford University Press.
- Silver, David, Hubert, Thomas, Schrittwieser, Julian, Antonoglou, Ioannis, Lai, Matthew, Guez, Arthur, Lanctot, Marc, Sifre, Laurent, Kumaran, Dharshan, Graepel, Thore, Lillicrap, Timothy, Simonyan, Karen, and Hassabis, Demis. 2018. "A general reinforcement-learning algorithm that masters chess, shogi, and Go through self-play." *Science* 362:1140–4.
- Singer, Daniel J. 2019. "Permissible epistemic tradeoffs." *Australasian Journal of Philosophy* 97:281–93.
- Smith, Angela. 2005. "Responsibility for attitudes: Activity and passivity in mental life." *Ethics* 115:236–71.

- Steglich-Petersen, Asbjørn. forthcoming. "An instrumentalist unification of zetetic and epistemic reasons." *Inquiry* forthcoming.
- Steglich-Petersen, Asbjørn and Skipper, Matthias. 2019. "An instrumentalist account of how to weigh epistemic and practical reasons for belief." *Mind* 129:1071–94.
- Steinberger, Florian. 2019. "Consequence and normative guidance." *Philosophy and Phenomenological Research* 98:306–28.
- Sylvan, Kurt. 2020. "An epistemic non-consequentialism." *Philosophical Review* 129:1–51.
- Thomsen, Frej Klem. 2011. "The art of the unseen: Three challenges for racial profiling." *Journal of Ethics* 15:89–117.
- Thorstad, David. 2021. "Inquiry and the epistemic." *Philosophical Studies* 178:2913–28.
- . 2022a. "There are no epistemic norms of inquiry." *Synthese* 200:410.
- . 2022b. "Two paradoxes of bounded rationality." *Philosophers' Imprint* 22:1–16.
- . 2023. "Against the newer evidentialists." *Philosophical Studies* 180:3511–3532.
- . forthcoming a. "The accuracy-coherence tradeoff in cognition." *British Journal for the Philosophy of Science* forthcoming.
- . forthcoming b. *Inquiry under bounds*. Oxford University Press.
- . forthcoming c. "Why bounded rationality (in epistemology)?" *Philosophy and Phenomenological Research* forthcoming.
- Treanor, Nick. 2014. "Trivial truths and the aim of inquiry." *Philosophy and Phenomenological Research* 89:552–559.
- . 2018. "Truth and epistemic value." *European Journal of Philosophy* 26:1057–1068.
- Tversky, Amos and Kahneman, Daniel. 1974. "Judgment under uncertainty: Heuristics and biases." *Science* 185:1124–31.

—. 1983. "Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment." *Psychological Review* 90:293–315.

Way, Jonathan. 2012. "Transmission and the wrong kind of reason." *Ethics* 122:489–515.

Whitcomb, Dennis. 2018. "Some epistemic roles for curiosity." In Ilhan Inan, Lani Watson, Dennis Whitcomb, and Safiye Yigit (eds.), *The moral psychology of curiosity*, 217–38. Rowman and Littlefield.

Wilard-Kyle, Christopher. forthcoming. "The knowledge norm for inquiry." *Journal of Philosophy* forthcoming.

Woodard, Elise. forthcoming. "Why double-check?" *Episteme* forthcoming.

Woodard, Elise and Flores, Carolina. forthcoming. "Epistemic vigilance: In defense of epistemic norms on evidence-gathering." *Philosophical Studies* forthcoming.