Rules for the Inquiring Mind—
A Unified Framework of Norms of Inquiry

Luis Rosa
Washington University in St. Louis

This is a preprint

Table of contents

§4.4 That we can now build a system of instrumental norms of inquiry on the basis of our theoretical constructs

Chapter 5

§5.1 That there are reasons to reject the Go Gather Norm
§5.2 Where we explore ways to salvage our framework
§5.3 Where we explore ways to salvage our framework (continued)
§5.4 Where we counter those objections to the Go Gather Norm

Chapter 6

§6.1 Where we also submit the Go Figure Norm to critical scrutiny
§6.2 That the 'should' of our norms of inquiry doesn't mean the same as 'must' does
§6.3 Where we start integrating different norms of inquiry into a unified system of norms
§6.4 That the Ignorance Norm of Inquiry, as presented above and justified in the way it was, leaves a number of important details implicit
§6.5 Where we also make those details of the Ignorance Norm of Inquiry explicit and we introduce yet another purported norm, namely, the Knowledge Norm of Inquiry

Chapter 7

§7.1 That the Knowledge Norm of Inquiry also leaves important details implicit
§7.2 That the Knowledge Norm of Inquiry entails the No False Presupposition Norm and the Anti-Dissonance Norm
§7.3 Where we become doubtful about including the Knowledge Norm of Inquiry into our system of norms
§7.4 That inquiry into Q is aimed at knowing Q
§7.5 That inquiry into Q is aimed at knowing Q (continued)

Chapter 8

§8.1 That our project wouldn't need to be completely abandoned were we to take the settling of a question Q to be more, or less, than knowing Q
§8.2 Where the Anti–Impossibility Norm is introduced to flag cases of misguided inquiry
§8.3 That the Anti-Impossibility Norm entails the No False Presupposition Norm, though not vice-versa
§8.4 That the No False Presupposition Norm may or may not flag the impropriety of inquiry into indeterminate questions, and that yet another norm of inquiry may lie strictly between the Anti-Impossibility Norm and the No False Presupposition Norm

Introduction

This book is about *inquiry*, and also about *questions*.

By 'inquiry' I mean a certain kind of *activity*—one that agents like you and me engage with more or less regularly. It is the activity of searching for the true answers to questions, of attempting to settle them. By 'questions' I mean a certain kind of *content*. Questions can be the contents of our mental states, as well as the contents of interrogative sentences and complements.

The investigation that follows delivers both, a number of theses about the *nature* of inquiry, and a system of *norms* of inquiry. It tackles not only philosophical issues regarding what inquiry is, but also issues about how it should and should not be. These two sets of issues are intertwined, of course, not at all independent of each other. Their interdependence is made evident at different parts of the book.

The views concerning the nature of inquiry defended here include the following:

(a) That knowing the answer to a question is a constitutive goal of inquiring into that question.
(b) That there are two general modes of inquiry, namely, information-gathering inquiry and armchair inquiry, occurring in tandem as they may.
(c) That progress in armchair inquiry hinges on the manipulation of *fine-grained information* (representational vehicles paired with content), without necessarily altering the inquirer's *coarse-grained information* (just content).
(d) That inquiry is always motivated by a combination of a doxastic and a volitional element—an interrogative attitude such as *doubt* or *uncertainty* towards a question, on the one hand, and a *desire* or *need* to know the answer to that question, on the other.

Some of the norms defended in the book are *proactive* norms, others are *prohibitive* norms. One of the proactive norms says that the inquirer should collect more information, given that she wants to settle some question and their information is not yet enough to settle that question. Another one says that the inquirer should see to it that they deduce a fine-grained answer of a certain type, given that they want to settle a question they already have enough information to settle. The prohibitive norms include the the following:

(e) That one shouldn't inquire into questions whose answers are *unknowable* to one.
(f) That one shouldn't occupy oneself with *pseudo-questions*.
(g) That one shouldn't inquire into questions one already knows the answer to.

These norms are understood as *instrumental* norms of inquiry. I specify the scope and the nature of their deontic operators, and I justify them using both, pre-theoretical considerations and the theoretical framework developed in the course of the investigation. That framework makes use of *models* of inquiry. A whole system of norms of inquiry is built on that basis.

The rough formulations of the target norms from above hide a number of important details. In particular, their explicit formulation features a *guise* parameter. A guise is a representational vehicle through which a question or a bit of information presents itself to an inquirer. It matters how a question is presented to an inquirer. They can inquire into a question as presented under one guise, without inquiring into it as presented under an alternative guise. They can be in doubt or uncertain about what the answer to the question is when it is presented under one guise, not when it is presented under a different guise. These important details have not been sufficiently explored in the recent literature on inquiry yet—the book fills that gap.

All in all, the book offers an overarching picture of what inquiry is and how it should be. I now proceed to summarize each of its chapters.

Chapter 1 does some conceptual groundwork. It characterizes inquiry as a certain type of goal-directed activity. It also makes two important distinctions: the distinction between *questions* and *interrogative sentences*, on the one hand, and the distinction between *information-gathering* and *armchair* inquiry, on the other.

Chapter 2 explicates the notions of *settling a question* and *that of information that is enough to settle a question*. A more or less standard construal of questions from formal semantics is adopted, where questions are taken to be the sets of their complete answers.

Chapter 3 explores a possible worlds model of inquiry, which represents successful information-gathering inquiry as the *elimination of previously uneliminated possibilities*. But that model fails to account for the value of *deduction* to inquiry. It is proposed that competent deduction changes the inquirer's *fine-grained* information, without changing their *coarse-grained* information.

Chapter 4 introduces a new model of inquiry which does capture the value of deduction to inquiry. The model vindicates a norm that directs inquirers to gather more information when they satisfy certain conditions (the Go Gather Norm) and a norm that directs inquirers to reason deductively from the information they already have when they satisfy alternative conditions (the Go Figure Norm).

Chapter 5 submits the Go Gather Norm to critical scrutiny. It responds to an objection that applies to that norm when the notion of *information that is enough to settle a question* is understood as suggested in Chapter 2 (the worry is that the Go Gather Norm becomes *too strong* when interpreted in that way).

Chapter 6 submits the Go Figure Norm to critical scrutiny, and objections to it are responded to. The Ignorance Norm of Inquiry—which says that one shouldn't inquire into *Q* when one already knows *Q*—is shown to be vindicated by the theoretical framework endorsed in the book. Nevertheless, it calls for a number of important amendments, in particular: (i) the norm is best understood as featuring ascriptions of *explicit* knowledge and, (ii) the norm is best understood as featuring a *guise* parameter. The chapter introduces yet another purported norm of inquiry, the Knowledge Norm of Inquiry, which states that one should inquire into *Q* only if one knows that *Q has a true maximal answer*.

Chapter 7 argues against and ultimately rejects the Knowledge Norm of Inquiry. The Knowledge Norm of Inquiry entails two other norms of inquiry that do seem to be true, however: the No False Presupposition Norm (that one shouldn't inquire into questions that rely on false presuppositions) and the Anti-Dissonance Norm (that one shouldn't inquire into questions one knows to admit of no true answers). The constitutive goal of inquiry into *Q* is here taken to be that of *knowing Q*. Objections to this view are presented and responded to.

Chapter 8 argues for the Anti-Impossibility Norm of Inquiry, which says that one shouldn't inquire into a question if it is impossible for one to know its true answer. This norm also entails the No False Presupposition and the Anti-Dissonance norms, though it does *not* entail the already rejected Knowledge Norm of Inquiry. Furthermore, the Anti-Impossibility Norm entails a norm against inquiry into indeterminate questions, or questions whose answers have indeterminate truth-values. The Go Gather Norm and the Go Figure Norm are shown to be consistent with the Anti-impossibility Norm, despite appearances to the contrary.

Chapter 9 makes the point that, even though the norms defended here are *instrumental* norms of inquiry, it is sometimes instrumentally beneficial for inquirers to violate them—but that this fact by itself doesn't refute those norms. It is also argued that, besides wanting or having a need to know *Q*, an agent who is inquiring into *Q* must at least be *uncertain* about *Q* (uncertain about which of the alternative answers to *Q* is the true one). Finally, it is argued that there is a clear sense in which it is possible for an inquirer to inquire into a question they already have the answer to, namely: they are in possession of a *coarse-grained* answer to their question, without being in possession of a *fine-grained* answer to that question as presented under a certain guise.

Chapter 10 tells norms of inquiry that impose epistemic requirements upon interrogative attitudes apart from the instrumental norms of inquiry advanced in the book. The system of norms that is built here countenances only the latter, not the former kind of norm. The rest of the chapter is occupied with fleshing out a norm of inquiry that tells us not to occupy ourselves with *pseudo-questions*.

Chapter 11 presents the main constituents of the system of instrumental norms of inquiry in a hierarchical order, dividing them into axioms and theorems, though somewhat informally. It sums up, then, the book's main outputs regarding the normativity of inquiry. It also recapitulates the grounds upon which that system was built and points at possible future extensions of the present investigation.

Chapter 1

§1.1 That inquiry is a goal-directed activity

Inquiry is something that minds do. It is a kind of activity that crucially involves *thinking*.

Examples of thinking include supposing that something is the case and drawing consequences from that supposition, trying to reconstruct past events, weighing the pros and cons of alternative hypotheses to explain the evidence, imagining an object to figure out what would happen to it under certain circumstances, jogging the memory, etc.

Thinking is constitutive of inquiry—but inquiry often involves more than just thinking. It may also involve actions or bodily movements such as opening doors, dialing numbers and walking towards specific sites. The process of thinking itself can be conceived of as involving bodily actions, of course, as when the inquirer writes a sentence down that expresses the supposition she's making and applies certain derivation rules to that sentence (thinking with a pencil), or when the inquirer mutters words to herself to prompt some memory recall (thinking with the mouth).[1]

Inquiry is not any old kind of thinking/thinking plus acting. It is an activity that is guided at finding the true answers to questions, or at settling them. The way we describe inquirers and their activities in natural language makes that much clear. An inquirer always inquires into *whether* something is the case, or *who* did so-and-so, or *when* this-and-that happened, or *why* it happened, etc. Interrogative complements such as 'where Waldo is' and 'who let the dogs out' express or have *questions* for their contents, which are the same as the contents expressed by interrogative sentences such as 'Where is Waldo?' and 'Who let the dogs out?' respectively.[2] We can always describe an inquirer or say what she is doing by using a declarative sentence that embeds one such interrogative complement.[3] For instance, we can use the declarative sentence 'Mary is inquiring into whether the universe started with a Big Bang' to

---

[1] Relatedly, see Clark and Chalmers (1998) for the view that various sorts of cognitive processes are realizable in external media. See also Bernecker (2014) for a critical take on debates about the hypothesis of an extended mind.

[2] This is not to say that interrogative complements *always* have questions for their contents, in *all* the linguistic contexts where they occur. See Belnap (1983) and Groenendijk and Stokhof (1984, Ch.2) for more on this.

[3] And so Dewey: 'Inquiry and questioning, up to a certain point, are synonymous terms' (1938, p. 105).

describe what Mary is doing. Where there is inquiry, there is an attempt to settle some question or other.

In a very simplified manner, we can reconstruct the activity of inquiry as follows: the inquirer raises herself a question and sets off to find an answer to it. Inquiry is a goal-directed activity, aimed at settling a question. Does that mean that the inquirer must not *already* have an answer to her question, in order to count as inquiring at all? Not necessarily—she might inquire into a question she already has the answer to. She might have that answer without realizing that she already has it, for example, or simply fail to be responsive to the fact that she already has the answer.

The issue is somewhat intricate and will occupy us again later (§9.5). For now, however, we can think of the inquirer as still curious or not yet satisfied, in that she is still looking for an answer to her question, or at least aiming to improve her epistemic standing with respect to one such answer (even in the latter case the subject's performance is still guided at settling a question).

This simplified reconstruction of the activity of inquiry is not supposed to imply that the inquirer formulates an interrogative sentence expressing the question she is inquiring into in her thought. Granted, it states that the inquirer 'raises herself a question'—but that is again some kind of reconstruction, rather than a literal description of the activity of inquiry.

Fido the dog has buried his bone in the backyard a few days ago, and now he is searching for that bone again. He uses his spatial memory to guide his search, but he still doesn't remember exactly where the bone is. We can describe Fido, if only with a bit of strain, as inquiring into *where the bone is*. That the dog doesn't token anything like an interrogative sentence in his thought doesn't itself prevent us from so describing him (not anymore than it prevents us from so describing a human being). The dog is also executing some kind of thinking activity, and we will interpret whatever representational vehicles we describe him as operating with—pictorial or imagetic as they may be—as having content as well.[4]

---

[4] What I write here concerning mental content is supposed to be compatible with anti-realism about representation and other forms of deflationism of representation talk—for example, Dennett's (1987) 'intentional stance' view, according to which a system counts as intentional if the assignment of mental states with content to it successfully predicts and explains its behavior. See also Lewis (1974) for a view according to which having propositional attitudes is a matter of being interpretable in certain ways. Such forms of anti-realism or deflationism do not forbid us from ascribing representational states or attitudes, and neither to they have to take them to be false, shall some paraphrasing mechanism of representation talk be available that describes things as they are without ontological commitment to representations. See Chapters 3 and 4 below for talk of representational vehicles, which will play a very important role here (footnote 9 of Chapter 3 has more references on deflationism about representation).

'Question' here also means a kind of *content*, as opposed to an interrogative sentence or complement. Questions are the contents of interrogatives, not the interrogatives themselves.[5] And questions can also constitute mental content—they can also be the contents of our thoughts.[6] Even if Fido is not literally posing a question to itself in the form of an interrogative like 'Where is the bone?', he is still related to the question of where the bone is, understood as content. Fido is still uncertain about *where the bone is* and is for that reason looking for it. The question of *where the bone is* constitutes the content of Fido's mental state.

The activity of inquiry always takes place because the inquirer is driven towards finding an answer to some question. The inquirer somehow wants or needs to know the answer to the target question. But that being so doesn't require her to formulate the question to herself using some interrogative.

All of that is compatible with saying, on the other hand, that an agent who is equipped with language has the opportunity to inquire into many more questions—and *more abstract* questions—than an agent who is not equipped with language.[7] It is difficult to conceive of the possibility, for example, that we could raise questions such as the following if we didn't have a faculty of language: (a) Did the universe have a beginning? (b) Which numbers satisfy the inequality $2x < 2$? (c) Are colors mind-independent properties of objects? With more language and its impressive combinatorial power comes more content for thought.

As it will be revealed below, getting more precise about the relationship between representational vehicles (including sentences) and content is key to addressing and hopefully solving some philosophical problems concerning inquiry.

§1.2 That inquiry is therefore not an attitude or a combination/series of attitudes

Inquiry is investigation—it is something that we *do* or engage with. It is not itself an intensional attitude of a specific kind, on a par with beliefs and desires, for example. And that is *not* because, in contrast to intensional attitudes, inquiry cannot last for long stretches of time (it can).

---

[5] And so van Fraassen (1980, pp. 137–138): 'A *question* is an abstract entity; it is expressed by an *interrogative* (a piece of language) in the same sense that a proposition is expressed by a declarative sentence'. See however §6.4 for the notion of a *fine-grained* question, which is made of content *and* a representational vehicle that carries that content.

[6] We are here following Friedman (2013) in thinking of questions also as *the contents of mental states*, over and above the fact that they are the semantic values of interrogatives. See also Carruthers (2018) for the claim that many non-human animals also have curiosity/questioning attitudes.

[7] For a similar point, but concerned with intensional attitudes such as belief and desire (rather than inquiry), see Dennett (2002, p. 559).

Many activities are extended through long stretches of time, and the activity of inquiry is one of them. Many such activities have to be put in *stand-by* mode, so to speak, waiting for their completion. Consider, for example, the activity of building a house. We say things like 'They are building a house', and that remains true for a long while. First they prepare the construction site, then they pour the foundation, then they put up the framing, install the plumbing, etc. Between each of these stages, even during some of these stages, the builders have to go home and get some rest and do other things. But all the while we say 'They are building a house', and truly so. The building of the house can take months. And so it is with inquiry—subjects can remain engaged in this activity for long periods of time, too.

So inquiry, just like intensional attitudes, can also last and be extended through long stretches of time. The reason why inquiry is not an intensional attitude, simply put, is this: that to inquire is to perform an activity, and to have an intensional attitude is not to perform an activity. Rather than being itself an intensional attitude, inquiry is a process that *takes* the subject from a state where he holds certain attitudes (the old state) to a state where he holds other attitudes (the new state). Ideally, the inquirer ends up with more knowledge on the topic inquired into at the end of the process, therefore with more beliefs on that topic at the end of the process. What was once an object of uncertainty or doubtfulness is now settled.

Inquiry constitutively involves the holding of intensional attitudes without being any of them. To inquire into a question *Q* is not the same as to hold an attitude of suspended judgment about *Q*, for example, or to be in doubt about *Q*.[8] Inquiry is rather a complex combination of thinking (reasoning, imagining, jogging the memory) and acting, all performed against the background of doxastic attitudes and other attitudes such as doubts, intentions and desires.

Inquiry doesn't even boil down to the inquirer's holding the attitudes that she does at different times in a chronological series. Suppose that, at the beginning of a process of inquiry, Dana starts off satisfying some open formula $f_1(x)$ that consists of a big conjunction of ascriptions of attitudes or their denials. For example, $f_1(x)$ could be: *x believes that humans are not aliens and x suspends judgment about whether there are aliens and x believes that there are other planets than Earth and x doesn't doubt that other planets than Earth are life-friendly and… and x wants to know whether there are aliens*. To say that Dana satisfies $f_1(x)$ is just to say that f_1(*Dana*), that is, that *Dana believes that humans are not aliens and etc.*

---

[8] See Friedman (2017) for the view that one inquires into *Q* if and only if one suspends judgment about *Q*.

As Dana comes to satisfy $f_1(x)$, she sets off to inquire into *whether aliens exist*: she starts reading about the size of the universe, collecting reports about alleged UFO sightings, hearing what chemists and biologists have to say about the necessary conditions for intelligent life on other planet, etc. As her process of inquiry develops, she transitions into new doxastic states. So at time t_1 she satisfies $f_1(x)$, a time t_2 she satisfies some other open formula $f_2(x)$, at t_3 she satisfies $f_3(x)$ and so on. Each of these formulas contains ascriptions of different doxastic attitudes or their denials.

Say that Dana resumes her investigation at t_n, where she satisfies $f_n(x)$. Hopefully, $f_n(x)$ features either *x believes that aliens exist* or *x believes that aliens do not exist* as one of its conjuncts, so that Dana has settled on an answer to her question at the end. And, hopefully, $f_n(x)$ does *not* feature *x wants to know whether there are aliens* as one of its conjuncts, so that Dana's isn't still thirsty for an answer at the end of her process of inquiry. Either way, what makes it the case that Dana satisfies $f_1(x)$, $f_2(x)$,…, $f_n(x)$ in this order is the fact that she has inquired into *whether aliens exist* during the time interval $[t_1, t_n]$.

Now we can remake the irreducibility point from above as follows. There could be another agent, say Jono, who satisfies exactly the same attitude-ascribing formulas $f_1(x)$, $f_2(x)$,…, $f_n(x)$ in the same chronological order, though Jono has *not* inquired into *whether aliens exist* during the time interval $[t_1, t_n]$. That is, both things are true:

(D) $f_1(Dana)$ & $f_2(Dana)$ &…& $f_n(Dana)$,
(J) $f_1(Jono)$ & $f_2(Jono)$ &…& $f_n(Jono)$.

And, yet, Dana but not Jono was inquiring into *whether aliens exist* during $[t_1, t_n]$. Just as (J) by itself doesn't entail that Jono was inquiring into *whether aliens exist* during the time interval $[t_1, t_n]$, so (D) does not entail that Dana was inquiring into *whether aliens exist* during $[t_1, t_n]$. Therefore, the fact that Dana was engaged in inquiry into that topic during that time interval doesn't just boil down to the fact that she held the attitudes ascribed by $f_1(x)$, $f_2(x)$,…, $f_n(x)$ at times t_1, t_2,…, t_n respectively.

What information is missing from $f_1(x)$ & $f_2(x)$ &…& $f_n(x)$ but is contained in an ascription of inquiry into the matter of alien life? *Something* will always be missing, no matter how rich and detailed we make each of the combinations of ascriptions of attitudes or their denials in each of the $f_i(x)$ of our sequence. One thing that is missing is a description of the subject's *thinking processes* that mediate the transitions from $f_1(x)$ to $f_2(x)$ and so on—for, again, thinking is an essential component of inquiring (it is the glue of inquiry).

Dana the inquirer has transitioned from one intensional state to the next *because* she was thinking about the issue of alien life between those transitions. For

example, say that from t_2 to t_3 Dana has lowered her credence that *there is life in another planet in the solar system* (other than Earth*)* because at t_2 Dana has learned that *there is no life in Mars* (one down), and she believes that *Mars is one of the planets in the solar system* at both times. This was achieved by her through reasoning, which in this case could be formalized as the lowering of credence in a disjunction (*either there is life in Venus or there is life in Mars or there is life in Saturn or…*) upon assigning lower credence to one of its disjuncts or even ruling it out completely (in this case, the disjunct is *there is life in Mars*).[9]

But Jono, we can assume, was transitioning from and to the very same kinds of intensional states *despite* the fact that her transitions were not mediated by her thinking processes concerning issue. For example, Jono might have lowered her credence that *there is life in another planet in the solar system (other than Earth)* from t_2 to t_3 because of some external interference in her neural circuitry.

It follows from this that, in order to fully assess Dana's inquiry, we also have to look at the thinking processes she went through in transitioning from t_1 to t_2, and then from there to t_3, etc. The fact that f_1(*Dana*) & f_2(*Dana*) &…& f_n(*Dana*) doesn't yet provide us with enough information to judge if Dana's inquiry has been properly conducted, for it doesn't yet inform us of the ways in which Dana has reasoned herself from one intensional state to the next.

Another *locus* of difference between Dana and Jono, of course, is action of a specific sort. Given that both Dana and Jono satisfy $f_1(x)$, they both want to know *whether aliens exist*. But only Dana *acts* on that desire and goes searching for relevant evidence because of it. Jono, we might suppose, is not so moved by that desire. Jono's changes of intensional states do not take place *because* she has actively searched for answers to the question of *whether alien exist*—she just happens to go through the same intensional updates as Dana's. Jono's performance during [t_1, t_n] is not in the game for being assessed as good or bad inquiry, for her performance wasn't an inquiring performance to begin with. And the intensional states that Jono has arrived at in each stage are similarly not in the game for being assessed as proper responses to the inquiring process that preceded them, liable to other forms of assessment as they may be (for example, we can still assess whether Jono's attitudes fit her evidence).

---

[9] See Staffel (2013, 2019) and also Dogramaci (2018) for the idea of reasoning with credences.

Considerations such as these make it clear that inquiry is not any an intensional attitude, or combination of intensional attitudes, or even a temporal series of combinations of intensional attitudes.[10]

§1.3 That some questions can be settled though armchair inquiry, others only through information-gathering inquiry

Inquirers will inquire into various sorts of questions. And we can distinguish questions under assorted dimensions, depending on what our theoretical goals are.

Consider, for example, the following contrasts: questions about the past, present or future *versus* questions about timeless truths; polar or yes/no questions (questions about *whether* something is the case) *versus* non-polar questions (questions about *who*…, *what*…, *when*…, *where*…, *how*…, *why*…);[11] the questions of philosophy *versus* the questions of the natural sciences; and so on. There won't be a one-to-one mapping between these distinctions, in the sense that every category from one of them corresponds to a category from the other.

To continue with the examples just introduced, not every question of philosophy is a polar question, and not every question of natural science is a non-polar question. The following non-polar question is a typical philosophical question: What conditions need to be satisfied in order for an act to be morally good? And the following polar question is a typical question from the natural sciences: Is the boiling point of Mercury above 300 °C? So a meta-philosopher who is occupied with the distinction between the questions of philosophy and the questions of the natural sciences, on the one hand, and the semanticist who is occupied with the distinction between polar and non-polar questions, on the other, will partition the general category of questions in different ways.

Just as there are different kinds of questions, so there are different ways to go about settling them. In particular, given the information that is already available to an inquirer at a given point in time, there are questions such that her being able to settle

---

[10] Does that mean that inquiry 'doesn't make a difference'? After all, Dana (the inquirer) and Jono (the non-inquirer) have achieved the same doxastic results going from t_1 to t_n… But we don't get to see the difference that inquiry makes to our intellectual lives by just comparing Dana's and Jono's *actual* doxastic developments. Rather, we have to compare Dana's actual doxastic development (the one that emerged out of her process of inquiry) with the one she *would have had* if she had not inquired into the target issue. If Dana hadn't been actively looking for answers and reasoning about the topic, it would most likely not be the case that $f_1(Dana)$ & $f_2(Dana)$ &…& $f_n(Dana)$.v

[11] Linguists, semanticists and logicians use different terminologies to mark this difference in particular—see for example Belnap and Steel (1976, §1.2), who distinguish what they call 'whether-questions' from what they call 'which-questions'. Ciardelli (2021, p. 31), among others, deploys the terminology used here, namely, 'polar questions' and 'non-polar questions'.

them at that time strictly depends on the possibility of her gathering *new information* from that time onwards, on the one hand, and questions that she is able to settle at that time *regardless* of whether she gathers new information from that time onwards, on the other.

Questions of the first type are such that the inquirer needs to acquire more information first in order to be in a position to settle them, and questions of the second type are such that the inquirer is already in a position to settle them—she does *not* need more information in order to be in that position. In the first kind of case, the inquirer needs to put herself in situations where she can see, hear, smell, touch and taste things, because doing these things might provide her with the target answer, and she wouldn't obtain that answer just by thinking and exploring the information she already has. In the second kind of case, however, the inquirer just needs to draw inferences from the information she already has, or more generally exploit that information through different forms of thought, in order to obtain the answer she is looking for. We could put it like this: there are questions that call for *information-gathering*, and questions that can be settled *from the armchair* (without gathering new information). Examples to illustrate the difference are now in order.

Suppose you are wondering *whether turtles lay eggs*. You have some information about turtles: you know that *they are reptiles*, that *they are cold-blooded*, that *they can stay underwater for a long time*, and you also know how they look like, that *they come in different sizes*, etc. But that information doesn't yet contain the answer to the question of *whether turtles lay eggs*. Say, for example, that even though you possess the information that *turtles are reptiles*, you do *not* posses the information that *all reptiles lay eggs*.

Moving forward, then, you consult a biology textbook and go to the section about reptiles. You read there that turtles do lay eggs indeed. You have then obtained an answer to your question *via* testimony. Your question here was such that, in order for you to settle it, you needed to acquire more information first—no amount of good reasoning from the information you already had would give you the answer you were looking for. Accordingly, you have done what you had to do for that purpose, namely, you have engaged in information-gathering inquiry.

In contrast, suppose you are wondering *what the powerset of the empty set* {} *is*. You know that *the powerset of a set is the set of all subsets of that set*. You know, for example, that *the powerset of* {1, 2} *is* {{}, {1}, {2}, {1, 2}}. You also know that *the empty set is a subset of every set*, and that *the empty set is a set without members*. From this, you can conclude that *the only subset of* {} *is* {} *itself*. So, via deduction, you reach the answer that *the powerset of the empty set is* {{}}.

In the last example your original question was not such that you had to gather new information first in order to be able to settle it. Accordingly, you have settled your question through sheer armchair inquiry—just by using your power of reasoning and the information you already had when you originally raised the question.

It might be suggested that the distinction between the questions that an inquirer is able to settle from the armchair and the questions that she will be able to settle only after acquiring new information neatly maps into the distinction between what is knowable *a priori* and what is knowably only *a posteriori* to that inquirer. The idea is that (a) the inquirer is able to settle a question *Q* through armchair inquiry just in case she can come to know the answer to *Q* in an *a priori* manner, and (b) the inquirer will be able to settle *Q* only after acquiring new information just in case the only way in which she can come to know the answer to *Q* is an *a posteriori* way.[12]

But given the typical way in which the notions of *a posteriori* and *a priori* knowability are understood, the suggested correlation does not really hold. Making trouble for (b) would be a bit more strenuous and controversial, so let us focus on (a).

Suppose the inquirer raises herself the question *Q* = *Was there a public holiday on April 15, 2023?* She remembers that *there was a public holiday three days ago*, and she also knows that *today is April 18, 2023.* Once she brings this information to bear on the issue, she is able to reason her way to an answer to *Q*, namely, that *there was a public holiday on April 15, 2023*. Before she did so, she was already in a position to settle *Q* from the armchair—the left-hand side of (a) was true of her.

But the proposition that *there was a public holiday on April 15, 2023* is not *a priori* knowable—so the right-hand side of (a) is not true of this case. For, presumably, one cannot come to know such a fact independently of empirical evidence or justification. Indeed, the fact that the inquirer was in a position to settle that question through armchair inquiry is due to the fact that she already had some such empirical evidence. If she didn't have such evidence, she wouldn't be in a position to settle her question from the comfort of the armchair. So the suggested correlation between armchair inquiry and the *a priori* does not hold.

Often a process of inquiry as a whole features both, episodes of information-gathering and armchair inquiry. We often put ourselves in a situation where we can gather relevant information first, and then we proceed to reason from the information thereby collected to a conclusion that answers our original question. For the collected information doesn't *itself* constitute an answer to the target question, though it is inferentially connected to an answer to it.

---

[12] For recent debates about the *a priori*, see the special volume of Casullo and Thurow (2013) and also Boghossian and Williamson (2020).

Conversely, sometimes we just process the information we already have in the armchair in order to decide where to search for relevant information next. And both of these sequences of events might be constitutive of larger, extended activities of inquiry.

Practical reasoning features in these cases centrally as a way of determining which actions the inquirer should perform in order to obtain the information she needs or wants. That includes the kinds of reasoning that go into designing experiments, selecting people to interrogate, which places to visit and scan the surrounding environment, etc. Practical reasoning, understood as the production of intentions through reasoning, is itself an armchair activity—but it leads to action that promotes the income of new information. And, where the new information is not yet the one we are looking for, but rather a stepping stone towards the information we are looking for, more reasoning is needed, be it theoretical or practical, or both.

Information-gathering and armchair inquiry are modes of inquiry—they are general *ways* in which we inquire into questions of various sorts. But since our inquiry into a given question often demands inquiry into other questions (the achievement of a goal requires the achievement of sub-goals), both ways of inquiring might be part of a larger, extended process of inquiry.

Depending on the question that the inquirer is trying to settle, armchair inquiry without new information is empty, and information-gathering without armchair inquiry is blind. It is true that every process of inquiry involves either information-gathering inquiry or armchair inquiry—but this is again not supposed to imply that a single process of inquiry cannot involve both modes of inquiry (that disjunction is not to be read as exclusive disjunction). We transition quite seamlessly between armchair inquiry and information-gathering inquiry in order to settle many of our questions.

Still, the questions that an inquirer is able to settle through armchair inquiry and those that she will be able to settle only after she gathers new information do *not* have any overlap. Reflection on these different possibilities reveals different kinds of problems. But before we proceed to explore these problems, we have to discuss yet another distinction to sharpen up the one that we have drawn in the present section.

§1.4 That further distinctions are needed

Upon closer inspection, the distinction between the questions that the inquirer is able to settle from the armchair and the ones that he will be able to settle only after he gathers new information needs to be refined. Just like an inquirer can either have or fail to have *information* that is enough to settle a given question, he can either have or fail to have the *cognitive skills* that are needed to settle that question.

So, for any inquirer and question *Q*, we can try to partition the space of possibilities thus:

(I^+S^+)	The inquirer has enough information to settle *Q*, and also the skills that are needed to settle *Q* on the basis of that information.
(I^+S^-)	The inquirer has enough information to settle *Q*, but not the skills that are nee to settle *Q* on the basis of that information.
(I^-)	The inquirer doesn't have enough information to settle *Q*.

We can draw the contrast between (I^+S^+) and (I^+S^-) as follows. In both cases, the inquirer has access to information on the basis of which *someone* could settle *Q*. But in (I^+S^-) the inquirer herself doesn't qualify as that someone, whereas in (I^+S^+) she does.

And now we can say that it is in all and only cases of type (I^+S^+) that the inquirer is able to settle the question through sheer armchair inquiry, without gathering any new information. In a situation of type (I^+S^+), the inquirer is *already* in a position to settle the question. In a situation of type (I^+S^-), the inquirer is *not* already in a position to settle the question (she is not actually able to settle that question through armchair inquiry).

In a situation of type (I^+S^-), the inquirer does have information that is enough to settle her question, but she doesn't yet have what it takes to *use* that information to flesh out an answer to that question. So when the inquirer is in a situation of type (I^+S^-), and assuming it is so much as *possible* for that inquirer to eventually settle her question, she doesn't strictly need more information in order to settle it. She can instead develop new cognitive skills that will eventually put her in a position to settle her question on the basis of the information she already has.[13] (Much of the value that we attach to learning *logic* stems from such a possibility, the hope being that learning logic will improve or enhance our reasoning skills).[14] Information-gathering is not a mandatory activity for her to successfully conclude her inquiry.

When thinking in terms of sets of questions, then, relative to any particular inquirer at any particular time, we can put part of the point through the equations:

The set of questions that the inquirer is able to settle through armchair inquiry at t =
The set of questions that the inquirer is already in a position to settle at t =

---

[13] The learning of new concepts, in some sense of that disappointingly obscure phrase, should also count as the acquisition of new cognitive skills. In many cases, it will simply consist of the acquisition of new *categorization* skills—the inquirer's *coarse-grained* information may remain the same before and after the acquisition of the new concept (see §2.2–2.3 for the notion of coarse-grained information), and what is new is just a way of partitioning the space of possibilities. Thanks to Dennis Whitcomb for discussion on this.

[14] See Manktelow (1999 Ch. 2–4) and also Stenning and van Lambalgen (2008, Ch. 5) and the references therein for experimental studies on this.

The set of questions such that the inquirer has both, enough information to settle them and the skills that are needed to settle them on the basis of that information at *t.*

When an inquirer already has the skills that are needed to settle a question on the basis of the information he already has, he can come to settle that question just by applying his actual cognitive skills to that body of information going forward (in the absence of any intervening factors such as heart attacks, etc.). He doesn't have to leave the armchair either to collect more information or to develop/sharpen his cognitive skills. Without worrying too much about redundancy, or about which skills from this list are built on the others, cognitive skills include: categorization and pattern-recognition skills, language-processing skills, the ability to visualize scenarios in a realistic manner, the ability to perform good inferences, to calculate, to retrieve information from memory, etc.

When an inquirer already has the skills that are needed for him to settle a question on the basis of the information he already has (I^+S^+), everything is ready for him to just go on and actively settle his question. In contrast, when he does *not* have those skills, more substantial transformations in his mental life are needed before he can settle his question—new ways of processing information must be implemented in his cognitive system. In the latter kind of situation (I^+S^-), it is *not* the case that the inquirer can come to settle his question just by applying his *actual* cognitive skills to the information possessed by him going forward (even in the absence of any intervening factors such as heart attacks, etc.).

We have just located cases where a subject is able to settle a question through armchair inquiry, or where the subject is already in a position to settle a question, in a more discriminating partition of the set of possibilities—one that also takes into account the inquirer's *cognitive skills*, and not only the *information* possessed by her.

Under this more fine-grained resolution of the modal space, we can now see that the questions that the inquirer is able to settle through armchair inquiry (I^+S^+) and the ones that she will be able to settle only after she acquires new information are not *exhaustive* of the set of questions.

That is because the set of questions that the subject *is not* able to settle through armchair inquiry is not identical to the set of questions that she will be able to settle only after acquiring new information. For some questions, relative to which (I^+S^-) is satisfied, the subject doesn't strictly need to acquire new information in order to be able to settle them, even though they are *not* questions that she is already in a position to settle (*ergo* they are not questions that she is able to settle through armchair inquiry, given only the information possessed by her).

Of course, it is hard to think of the possibility of a subject's developing new cognitive skills without thereby acquiring *some* information along the way. If the development of new skills must involve the acquisition of new information,[15] then there will be both cases of type (I^+S^-) and cases of type (I^-) such that the inquirer will be able to settle her question only after she first acquires new information. (Notice that there will also be questions satisfying (I^+S^-) and (I^-) such that it is impossible for the inquirer to settle them).

But when we talk about a question that the inquirer will be able to settle only after acquiring new information (going forward), what we mean to talk about is a question that the inquirer does not yet have enough information to settle. That is, we mean to talk about a question such that (I^-) is satisfied for it, relative to that inquirer at that time. That is why we say that the inquirer *strictly needs* to acquire more information in order to be able to settle her question—information, that is, such that, when added to the information she already has, will provide her with a total body of information that *is* enough to settle her question. It is *that* kind of information that she needs to gather. Only then will she be in a position to settle her question.

Accordingly, regarding cases of type (I^+S^-), when we deny that the inquirer will be able to settle her question only after she acquires new information going forward, we are not denying that the inquirer will be able settle her question only if she acquires *some kind* of new information going forward, including collateral information that comes with the acquisition of new cognitive skills.

What we mean to deny here, that is, is that the inquirer will be able to settle her question only if she acquires new information which *will transform her total information set into one that contains information that is enough to settle her question* (for she *already has* information that is enough to settle her question). In some cases of type (I^+S^-), the inquirer can come to be in a position to settle her question even *without* acquiring information with that specific property, or *regardless* of whether she acquires that kind of information (she can come to be in that position by learning new cognitive skills, plus whichever collateral information this learning will provide her with).

§1.5 That we now have a sharp and exhaustive distinction between questions

---

[15] That would be the case if a cognitive skill already *is* some kind of propositional knowledge. See Pavese and Beddor (forthcoming) for the view that the skills involved in skillful and intentional action are propositional knowledge states—though as far as I can see their view does not extend to all the skills that I am calling cognitive skills here (for example, memory retrieval skills).

Even though there is a sharp distinction between the questions that an inquirer can settle through armchair inquiry and the questions that he can settle only after collecting new information, that distinction is not yet *exhaustive* of the total set of questions.

The distinction is *sharp* because there is no question such that one already has information that is enough to settle it and does not have information that is enough to settle it at the same time. In the way we are using these terms of art here, the former is a necessary condition for the truth of 'one is able to settle the question through armchair inquiry' (relative to the information possessed at a time *t*), and the latter is a necessary condition for the truth of 'one will be able to settle the question only after acquiring new information' (relative to the information possessed at *t*).

The distinction is *not exhaustive*, however. And one way to see that is by noticing that the questions that the inquirer is able to settle through armchair inquiry are the questions that he is already in a position to settle, but that the questions that the inquirer will be able to settle only after he acquires new information is not simply the complement of the questions that he is already in a position to settle. The questions that he is *not* already in a position to settle includes two different sorts of questions: (I^-) the ones such that he doesn't already have enough information to settle them, and (I^+S^-) the ones such that, even though he already has enough information to settle them, he lacks the skills that are needed to settle them.

Now it is time to try a completely universal statement about inquiry and questions: for any inquirer at any time, every question falls into exactly one of the three categories that we have just carved out. Is this true?

Take any arbitrary question *Q*. Regarding a particular inquirer (at a particular time), then, now we ask: Does the inquirer have information that is enough to settle *Q*? If not, then *Q* is a question of type (I^-) for that inquirer. If yes, then *either* the inquirer has the cognitive skills that are needed to settle *Q* on the basis of the information possessed by her, or she doesn't. If the former, then *Q* is a question of type (I^+S^+) for that inquirer, and she is able to settle it through armchair inquiry. If the latter, then *Q* is a question of type (I^+S^-) for that inquirer, and it is not the case that she is able to settle it through armchair inquiry.

Since we have covered all the possibilities, it seems that any arbitrary question must indeed fall into one of our three categories—(I^+S^+), (I^+S^-) or (I^-)—for any inquirer at any particular time and possible world.

The following consideration, however, might make us doubtful about whether any arbitrary question falls into *exactly* one, and not more than one, of those three categories, for any inquirer at any time. It might happen that the inquirer's total body of information contains two subparts such that *either* of them is by itself enough to settle *Q*, and the inquirer has the cognitive skills that are needed to settle *Q* on the basis of

one of those subparts, but lacks the cognitive skills that are needed to settle *Q* on the basis of the other subpart. So doesn't *Q* fall into *both* categories, (I^+S^+) and (I^+S^-), for that inquirer at that particular time?

Let us for the moment think of the total information possessed by the inquirer as a set of propositions. So the kind of situation envisioned by the troubling consideration is one where the total information possessed by the inquirer at *t* is some set {*p*, *q*, *r*, *s*}, where both {*p*, *q*} and {*r*, *s*} are by themselves enough to settle some question *Q*, but at *t* the inquirer only has the skills that are needed to settle *Q* on the basis of {*p*, *q*}, not the skills that are needed to settle it on the basis of {*r*, *s*}.

When the problem is framed in this way, we can explain why *Q* does *not* fall into both categories (I^+S^+) and (I^+S^-) as follows. The possibility of type (I^+S^+) was phrased thus: 'The inquirer has enough information to settle *Q*, and also the skills that are needed to settle *Q* on the basis of that information'. Under the present construal of bodies of information, this should be read as saying that the inquirer has *some* information that is enough to settle *Q*—that is, some subset of the total set of propositions that constitutes the total body of information possessed by her—such that she has the cognitive skills that are needed to settle *Q* on the basis of *that* information (that same subset).

In a possibility of type (I^+S^-), by contrast, there won't be *any* subset of the total information set that it is both, enough to settle *Q* and the inquirer has the skills that are needed to settle *Q* on the basis of it—even though *there is* some subset of her total information set that is enough to settle *Q*.

As a result, it is not the case that the apparently troubling type of case falls into *both* categories, (I^+S^+) and (I^+S^-), for the envisioned inquirer and time. The target question *Q* belongs to the set of questions for which (I^+S^+) applies, and *only* to that set. It is therefore a question such that the inquirer is able to settle it through armchair inquiry, or she is already in a position to settle it.

Of course, if the inquirer tries to settle *Q* using armchair methods she might 'bark up the wrong tree' and explore the wrong subset to squeeze out an answer to *Q* from (the subset of her total information set on the basis of which she does *not* have the skills to settle *Q*). Still, however, she has everything that she needs to actually settle *Q*: some *other* subset of her total information set and the skills that are needed to settle *Q* on the basis of *that* piece of information. She doesn't strictly need either new information or new cognitive skills in order to settle her question—she just needs to make a better choice as to which chunks of her total information set to explore when attempting to do that.

Similar observations would apply to a version of the worry framed in terms of information 'fragments', understood as ways of enabling access to different bits of

information possessed by the inquirer.[16] A question *Q* will fall into category (I^+S^+) but not (I^+S^-) for the inquirer when there is *a* fragment relative to which the subject has access to information that is enough to settle *Q* and she has the skills that are needed to settle *Q* on the basis of that information.

It seems, then, that our universal statement about inquiry and questions is indeed true: every question *Q* falls into one and only one of the categories (I^+S^+), (I^+S^-) or (I^-), for any inquirer at any time.

Once we accept that universal statement, we might feel tempted to look for even stronger theses relating types of questions to inquiry. In particular, we might feel tempted to ask whether there are questions such that it is *impossible* for an inquirer to settle—or maybe it is even impossible for *any* inquirer to settle. Our taxonomy so far suggests that there are two ways in which that could be so: it might either be impossible for the inquirer(s) to have enough information to settle those questions, or impossible for the inquirer(s) to have the skills that are needed to settle them. If either of these holds, then there are *unsettable* or *unanswerable* questions, though that does not mean that *there are no* answers to those questions. The issue will recur throughout this investigation (see §8.2 for a norm against inquiry into unsettable questions).

Now we need to get more precise about some of the notions that we have been deploying—crucially, the notions of *settling a question* and that of *information that is enough to settle a question*.

[16] See Elga and Rayo (2021) for this notion of fragmentation.

Chapter 2

§2.1 That settling a question is a form of success

The appropriate way for us to understand what it is for an inquirer *to settle* a given question must be such as to make her settling of the question a form of epistemic success or accomplishment. To see why, let us look back at the roles that the notion of settling a question has played in this investigation so far.

We have distinguished between the questions that an inquirer is able to settle through armchair inquiry and the questions that she will be able to settle only after she acquires new information. That is a difference between the questions that the inquirer is already in a position to settle and the questions that she will be a position to settle only after the income of new information. But, of course, in a relatively useless sense of 'settle a question', an inquirer will already be in a position to settle pretty much any question she happens to entertain, thus depriving the distinction of its theoretical value.

For example, suppose we take it that an inquirer settles a question *Q* when, after being prompted to find an answer to *Q*, she comes to believe one of the answers to *Q* (no matter how she does it). Now suppose you are inquiring into the question of *whether the first color movie was produced in Hollywood*. Say you have a hunch that the answer is *yes*—but that is only a hunch of yours. Despite the fact that you didn't yet weigh the available evidence pro and con that answer to that question, or hear what historians have to say about it, you are so moved by your strong hunch that your thereby come to believe that *the first color movie was produced in Hollywood*. Since that is indeed one of the question's answers, you count as settling it by the present standard.

So, before all of this happened, as you raised yourself the question, you were already in a position to settle it—because you were already in a position to simply come to believe one of that question's answers. For our present purposes, however, such a verdict is completely spurious. We have lost the opportunity to draw an epistemologically relevant contrast between your standing with respect to the question of *whether the first color movie was produced in Hollywood* (a question, we may assume, you were *not* in a position to settle) and any other question that you *really* were in a position to settle.

That notion of settling a question is just not the one that is fit for the job of defining what it is for one to be in a position to settle a question in the way we did. It doesn't square well with the stipulation that one counts as being in a position to settle a

question *Q* just in case one already has information that is enough to settle *Q*, as well as the cognitive skills that are needed for one to settle *Q* on the basis of that information. In order for this to come out true, settling a question *Q* must not simply consist of believing one of *Q*'s answers in any old way.

What kind of epistemic success must we require in order for one to count as settling a question, then? One idea is that settling a question consists of coming to *know* the question's true answer after being prompted to find an answer to it.[17]

Since knowledge requires truth, an inquirer will never count as settling a question when the answer she has arrived at is false. Suppose detective Constanza is inquiring into *who killed Mr. Jones*. Constanza has narrowed her suspects down to Anton, Betty and Carmen. Constanza later finds out that, around the same time of Mr. Jones' murder, Betty and Carmen were seen together in the same car about 1.5 miles away from Mr. Jones' apartment. Furthermore, Constanza believes on good evidence that *Mr. Jones was murdered in his own apartment*—there are signs all over indicating that the killing happened there. Good reasoning on the basis of that information leads Constanza to conclude that *it was Anton who killed Mr. Jones*.

So far, it seems that Constanza has settled the question of *who killed Mr. Jones*. What Constanza doesn't know, however, is that Mr. Jones was not murdered at his own apartment, all the appearances to the contrary notwithstanding. In fact, it wasn't Anton who killed Mr. Jones, but rather Betty and Carmen, who killed him in their car and brought his dead body to his apartment shortly afterwards. Furthermore, they have planted evidence to suggest that Mr. Jones was murdered inside the apartment.

According to the proposal under consideration, Constanza did not settle the question of who killed Mr. Jones. Constanza's conclusion was that it was Anton who killed Mr. Jones. But that is false, and therefore not knowledge. It might *seem* to Constanza as if he has settled his question—but here as elsewhere appearances are misleading.

At this point, however, we don't strictly need to decide whether settling a question consists of knowing the question's true answer (upon being prompted to find an answer to it). The notion of settling a question can play the theoretical role that we want it to play up to this point—namely, that of allowing us to draw the distinctions referred to above—even if we require *less* than that for its correct application. It will be enough for now to say that the inquirer counts as settling a question only if she comes to believe an answer to her question by means of a reliable belief-forming process (say, she believes the answer on the basis of good reasoning from the information

---

[17] See Kelp (2021) for discussion and comparison to other criteria for one to count as settling a question.

possessed by her, or as a result of performing other kinds of reliable thought processes that deploy/explore that information, if we don't want to call them reasoning processes).

So you do *not* count as settling the question of *whether the first color movie was produced in Hollywood* in the example discussed above, for you did not come to believe the *yes* answer to that question as a result of any reliable thinking process. We can still say that you were not in a position to settle that question before you went on and came to form your hunch-based belief in one of that question's answers.

We embrace, then, that minimal condition for the settling of a question, though we will strengthen it by imposing the knowledge requirement referred to above later (§8.1). To believe one of the question's answers by means of a reliable belief-forming process is already enough of an accomplishment to rule out many ways of reaching an answer to a question as candidates for settling a question.

To settle a question is a form of epistemic success, minimally, reliably formed belief in one of the question's answers. When the inquirer comes to believe an answer to her question through unreliable thinking processes such as wishful thinking, fallacious reasoning, confusion between memory and confabulation, etc., she hasn't really settled her question, thinking that she did as she may. And, if the inquirer's cognitive skills are not enough to allow her to believe one of the question's answers by executing a reliable thinking process, then that inquirer is not yet in a position to settle that question, *ergo* it is not the case that she is able to settle that question through armchair inquiry, even though she already has information that is enough to settle it. (See §5.2 for the distinction between the literal use of 'settle a question' and the use of that phrase in loose talk).

There are many ways of explicating the notion of reliability in the literature, and the task is beset by problems of various sorts.[18] Common to all of those explications is an attempt to precisify the vaguely put idea that a reliable belief-forming process is a belief-forming process that is conducive to truth. That notion is presupposed/relied on for many other theoretical purposes than the ones that currently occupy us. This is unsurprising, seeing as many epistemic forms of accomplishment (including knowledge) require the reliability of the belief-forming process.

Here we won't try and flesh out a theory or explication of reliability. Whatever more exactly makes it the case that a belief-forming process is a reliable or a truth-conducive one, an inquirer counts as settling her question only if she comes to believe one of the question's answers through a reliable belief-forming process.

---

[18] See for example Feldman (1985). See also Grundmann (2022) for problems to do with the notion of *conditional* reliability, and a proposal about how to set it straight.

§2.2 Where the notion of *information that is enough to settle a question* is further explicated

Little by little we start gaining a better understanding of our inquiry-related taxonomy of questions from Chapter 1.

We noted initially that there are questions that an inquirer is able to settle through armchair inquiry, just by thinking/reasoning from the information that he already has, and questions that he will be able to settle only after acquiring new information. The former questions are such that the inquirer is already in a position to settle them. That the inquirer is already in a position to settle them means that he has both, information that is enough to settle them and the cognitive skills that are needed for him to settle them on the basis of that information. And now we have clarified what is supposed to go under the verb 'settle' in these constructions: to settle a question is, perhaps among other things, to execute a reliable thought process that outputs a belief in one of the question's answers.

Now some clarifications about the expressions 'information' and 'information that is enough to settle a question' are called for. This is the task that we take up in the present section and the next one.

Information that is enough to settle a question *Q*, remember, is what an inquirer *lacks* when it is true of him that he will be able settle *Q* only after he acquires new information. The rough idea is as follows. We all have access to some information. All in all, an inquirer's total body of information contains the answers to some questions, but not to others. In the latter sort of case, there is a gap in the inquirer's total body of information. To fill that gap, more information has to come in, so that previously open questions become closed. And, when the inquirer's information already contains the answer to a given question, there is no gap for her to fill (she can find the answer within herself, so to speak).

A bit trivially, we can already say that information that is enough to settle *Q* is at least information such that there is a reliable thinking process that deploys that information to output a belief in one of *Q*'s answers.

But that is uninformative, because it simply unpacks one of our necessary conditions for the correct application of 'settles a question' from the previous section. And we cannot strengthen it by simply *identifying* what is referred to on the left-hand side to what is referred to on the right-hand side (by substituting 'is' for 'is at least').

In practice, thinking of a reliable thinking process of the relevant type—or of how one could go from one's total body of information to one of the question's answers in a reliably manner—is a good guide or heuristic to decide whether that body of information is enough to settle that question. But, on account of the fact that settling a question may involve more than just reaching one of the question's answers through a

reliable belief-forming process that deploys the information possessed by the inquirer, the existence of such a process won't *always* tell us that that body of information is enough to settle that question. Various sorts of cases constitute potential counter-instances to the hypothesized identity, including cases where the inquirer reliably extrapolates the data available to her via some form of abductive or inductive reasoning (fallible but reliable reasoning processes—more on this in §§5.1–5.3).

For example, there may be a reliable process of reasoning that leads the inquirer from a body information that registers the fact that *all taxi cabs that have been observed in NYC so far are yellow* (observed by the inquirer within a given time-interval) to a belief in the generalisation that *all taxi cabs in NYC are yellow*. But perhaps the statistical data possessed by the inquirer doesn't yet contain the answer to/doesn't yet close the question of whether *all taxi cabs in NYC are yellow*. That question may still be an open question, as far as the statistical data in the inquirer's total body of information goes. The statistical data doesn't yet *settle* the question. (Compare the question of *whether it is probable that all taxi cabs in NYC are yellow*—the latter may be closed, depending on the meaning of 'probable', even if the question of *whether all taxi cabs in NYC are yellow* is still open).

Furthermore, a reliable thinking process that outputs a belief in one of the question's answer on the basis of a given body of information might be *more* or *less* reliable, whereas the question of whether it is enough information to settle the question is an on/off or a binary matter. Either the information is enough to settle the question or it isn't, with no shades of gray in between.

We can take inspiration from the literature on the semantics of questions to start making our notion of *information that is enough to settle a question* more precise and tractable, subject to idealizations as it may be.[19] To use a popular approach in this field, we think of a question as a set of answers.[20] Just like a question is not simply an interrogative sentence, but rather its content, the answer to a question is not simply a declarative sentence, but rather its content. Depending on what our theoretical preferences are, this amounts to saying that a question is a set of *propositions*. In the canonical framework, the answers to a question are sets of possible worlds or, equivalently, functions from possible worlds into truth-values (intensions). If we are happy to think of propositions as sets of possible worlds, then we will be happy to say that questions are sets of propositions.[21]

---

[19] See Roelofsen (2019) for an overview.

[20] Variations of this general approach can be found in Hamblin (1973), Karttunen (1977), Groenendijk and Stokhof (1984), Ciardelli et al. (2019).

[21] Lewis (1973), Kratzer (1977) and Stalnaker (1984), among others, adopt this view of propositions.

Otherwise, we simply say that the answers to a question are sets of possible worlds—or, as I will often put it here, they are bits of *coarse-grained information* (see below for more on the coarseness of grain of bits of information, where the alternative notion of *fine-grained information* is also introduced).

The latter is the theoretical option that we will adopt here. We will remain neutral on the issue regarding the nature of propositions, and we won't be committed to the view that propositions are sets of possible worlds or functions from possible worlds to truth-values. Instead of saying that *p* is the proposition expressed by '*p*', we simply say that the sentence '*p*' *carries* the coarse-grained information that *p*, which is the set of possible worlds where '*p*' is true (where '*p*' has the meaning that it has in *our* mouths in the context we utter this whole sentence). The answer to a question is a bit of coarse-grained information, namely, a set of possible worlds. The question is a set of such bits of coarse-grained information—those that constitute answers to that question.

There is more than one way of construing questions as sets of answers within this general framework of possible worlds semantics. But for our present purposes we will employ one construct in particular, which stems from work in inquisitive semantics, namely, that of a question as the downward closed set of its maximal answers (notion explained below).[22] Our claims about what it takes for a body of information to be enough to settle a question are translatable into other variations of the same general framework. The constructs of inquisitive semantics just happen to suit our purposes in a more straightforward manner. We illustrate its workings by looking at so-called *polar* or *yes*/*no* questions in the next section.

§2.3 Where the notion of *information that is enough to settle a question* is further explicated (continued)

Let *p* be the coarse-grained information carried by the sentence '*p*', and ¬*p* the coarse-grained information carried by the sentence 'it is not the case that *p*'. Then the question of *whether p* (a polar question), or *?p*, will be represented as a set containing at least both, *p* and ¬*p* as its members.

For example, the question of *whether there is going to be a Third World War* contains at least both of the following bits of coarse-grained information: that *there is going to be a Third World War* and that *there isn't going to be a Third World War*, understood respectively as the set of possible worlds where 'There is going to be a Third World War' is true (as uttered from our context) and the set of possible worlds where that sentence is false.

---

[22] See Ciardelli et al. (2019) for the use of such a construct in inquisitive semantics (which the authors call 'issues').

Both *p* and ¬*p* respectively are *complete* answers to the question of *whether p*. But they are not the *only* complete answers to it, since any other answer that *entails* either *p* or ¬*p* also counts as a complete answer to that question.

For example, that *the third World War will be ignited by an economical crisis* entails that *there is going to be a Third World War*, and so the former is also a member of the question of *whether there is going to be a Third World War*, as much as the former is. In this framework, entailment is captured by the subset-relation. That is, the information that *a* entails the information that *p* when *a* is a subset of *p*, or $a \subseteq p$ (every possible world where '*a*' is true is also a world where '*p*' is true).

That *a* also counts as a complete answer to *?p* when $a \subseteq p$ makes perfect sense, since in this case *a* is *at least as informative* as *p* is. If *p* is enough to close a given question, then *a* is enough to close it, too. So the question of *whether p*, or *?p*, will contain not only *p* and ¬*p*, but also any other *x* such that either $x \subseteq p$ or $x \subseteq \neg p$.

In other words, the set ?*p* is *downward closed*, starting from the maximal elements *p* and ¬*p*, called its *maximal answers*. Every maximal answer to a question is a complete answer to it, but not every complete answer to it is a maximal answer to it. Quite generally, a maximal answer *x* to a question *Q* is a complete answer to *Q* such that there is no other complete answer *y* to *Q* with $x \subset y$.

Let us use a toy model and suppose that the total set of worlds that is being used to interpret our language features four possibilities, w_1, w_2, w_3 and w_4. Suppose further that '*p*' is true in w_1 and w_2, but false in w_3 and w_4. So $p = \{w_1, w_2\}$ and $\neg p = \{w_3, w_4\}$. The space of possibilities is then divided thus:

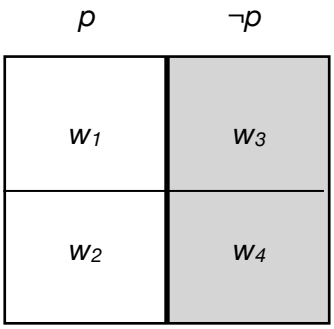

So the question of *whether p* here is $?p = \{\{w_1, w_2\}, \{w_3, w_4\}, \{w_1\}, \{w_2\}, \{w_3\}, \{w_4\}, \{\}\}$.

That is, it contains $p = \{w_1, w_2\}$, $\neg p = \{w_3, w_4\}$ and all the subsets of each, down to the empty set {}. Now not only $p = \{w_1, w_2\}$ is a complete answer to *?p*, but also, say, $\{w_1\}$, since $\{w_1\} \subseteq p$ (entailment relation).

We can put it like this: if the fact that the actual world lies on the left side of our diagram is enough to answer the question, then the fact that the actual world is w_1 is

also enough to answer that question, seeing as w_1 lies on the left side of the diagram. Similarly, not only $\neg p = \{w_3, w_4\}$ is a complete answer to *?p*, but also $\{w_3\} \subseteq \neg p$.

A complete answer to a question is any bit of coarse-grained information that is a *member* of that question (let us bracket questions about the empty set/what it means for now). Among the complete answers to a question, again, there are always the ones that are the *maximal* answers to it, in this case *p* and ¬*p* (we might think of them as the sets that 'start off' the process of getting the whole question to be downward closed).

In addition to complete answers, there are also *partial* answers. In this framework, partial answers are answers that, even though they eliminate some of the complete answers to a question, they do not entail any complete answer to it.

Let us continue with the example from above and imagine that *?p* is the question of *whether there is going to be a Third World War*. In w_1 and w_2 a Third World War indeed takes place, but not in w_3 and w_4. Suppose, however, that whereas the Third World War starts in Europe in w_1, it starts in the Middle East in w_2. So the bit of coarse-grained information that *the Third World War will start in the Middle East* is here the singleton set $\{w_2\}$.

Now the answer that *the Third World War (if there is one) won't start in the Middle East* is a *partial* answer to the target question, seeing as it only eliminates w_2 as a candidate for being the actual world, and it only tells us that the actual world is one of those in the set $\{w_1, w_3, w_4\}$. That *the Third World War (if there is one) won't start in the Middle East* is not a complete answer to the question of *whether there is going to be a Third World War*. But it is not an idle answer either—after all, it rules out at least one possibility where there is a Third World War.

Given one of its complete answers, a question is *closed*. But given only a partial answer to it, the question still remains *open*. Now we are much closer to explicating the notion of *information that is enough to settle a question*.

We have been saying that an inquirer has/possesses information, and that her information is enough to settle some questions but not others. In particular, an inquirer will also be in possession of coarse-grained information, which is some set of possible worlds.[23] This will be so even if the (unqualified) notion of information that is at play when we say that the inquirer possesses information is a more fine-grained one.

Say, for example, that we are thinking of the information possessed by the inquirer as a set of propositions relative to which the inquirer satisfies some epistemic

---

[23] This kind of model of the information possessed by an agent has its roots in Carnap and Bar-Hillel (1952), who flesh out ways of measuring the amount of information carried by sentences of a particular language. Hintikka (1962) also uses it to give truth conditions to ascriptions of belief and knowledge, understood either as ascriptions of implicit belief and knowledge or, alternatively, as ascriptions of belief and knowledge to ideal agents.

condition—the propositions that are known by the subject, or the ones that belong to her total body of evidence—and that such propositions are more fined-grained than sets of possible worlds.

A proposition might be, for example, a tuple whose members are the semantic values/intensions of the expressions of the sentence that expresses that proposition (a structured proposition).[24] Even though propositions themselves are not bodies of coarse-grained information (sets of possible worlds), the inquirer will *also* be in possession of coarse-grained information—say, the set of possibilities where every proposition that is known by her is true, or the set of possibilities where every proposition that is part of her evidence is true—in virtue of possessing fine-grained information. Those would be the set of possible worlds that are compatible with what the inquirer knows and the set of possible worlds that are compatible with her evidence, respectively.

We could also put it like this: the coarse-grained information possessed by an inquirer will be the set of worlds where the fine-grained information possessed by her is true. Fine-grained information is more intimately connected with *how the subject represents* the world as being or what her *representational states* are. We say that the subject has or possesses certain coarse-grained information when her representational states carry that information—but in a sense her coarse-grained information extrapolates her fine-grained information (more on this in §3.4).

A subject's representational states might carry certain bits of coarse-grained information without her being aware of what that coarse-grained information is, without this being explicit or even accessible to her. There are no guarantees that a subject will know what bits of coarse-grained information are possessed by her.

Some questions will be closed, others open, relative to the coarse-grained information possessed by the inquirer. The coarse-grained information *i* possessed by the inquirer is some set of possible worlds—say, again, that *i* is determined by the propositions that are know by her, or by the propositions that count as part of her total evidence.[25] And we are representing questions as sets containing bits of coarse-grained information.

So a question *Q* is *closed* relative to the coarse-grained information *i* possessed by the inquirer just in case *i* is one of the members of *Q*, or $i \in Q$. Otherwise *Q* is still *open* relative to the information *i* possessed by her. In the former case, then, the inquirer has information that is enough to settle *Q*, but in the latter case she doesn't. An

---

[24] See for example Lewis (1972) and Cresswell (1985) for variations of this view.

[25] That is, *i* is a set of possible worlds such that, for every possible world *w*, $w \in i$ if and only if every proposition that is known by the subject/is part of the subject's evidence is true at *w*.

that is our explication of what it is for an inquirer to have information that is enough to settle a question.

If we were to formalize a question as the set of its *maximal* complete answers, without being downward closed, we would say: an inquirer's information *i* is enough to settle *Q* just in case *i* is a *subset* of one of the members of *Q*. Since our questions are downward closed, however, we can just say: an inquirer's information *i* is enough to settle *Q* just in case *i* is a member of *Q*.[26] What this means in effect is that the agent is in possession of a *complete answer* to the question, where the answer is understood as a bit of coarse-grained information.

If, for example, all the possible worlds compatible with the information available to you are such that *this book was written by a human being*, then you have an answer to the question of *whether this book was written by an artificial intelligence system* (we are assuming you possess the information that *no human being is an artificial intelligence system*). But that you are in possession of an answer to this question in this sense doesn't entail, again, that this is explicit to you or that you know that you have that answer (consider for example the time when you were reading the previous paragraph).

To use the toy model from above, consider again the question of *whether p*, namely, *?p* = {{w_1, w_2}, {w_3, w_4}, {w_1}, {w_2}, {w_3}, {w_4}, {}}. Suppose the coarse-grained information possessed by the inquirer is {w_1, w_2}. Under our supposition, then, the inquirer possesses information that is enough to settle *?p*. In fact, we were supposing that the inquirer possesses the information that *p* (though she might not know this), and of course *p* is a member of *?p*. The question *?p* is closed relative to *p*, or *p* is a complete answer to *?p*.

Things would be different if the information possessed by the inquirer were, say, *q* = {w_2, w_3}. For notice that *q* is *not* a member of *?p*. The question *?p* is still open relative to *q*, as *q* is only a partial answer to *?p*. It rules out w_1 and w_4 as candidates for being the actual world, while leaving it open whether it is w_2 or rather w_3 that is the actual world. But since w_2 is a *p*-world and w_3 is a ¬*p*-world, the information possessed by the inquirer in this case is not yet enough to settle the question of *whether p*.

To sum up, an inquirer will always possess some information, which might be thought for example as the propositions (even fine-grained propositions) that are known by her to be the case, or the ones that are part of her evidence. To that information corresponds a body of coarse-grained information—namely, the set of possibilities where all the target propositions are true.

---

[26] Ciardelli et al. (2019) call that 'issue resolution'. See Groenendijk and Stokhof (1984) for the construal of questions as the sets of their maximal answers, but not all their complete answers.

The coarse-grained information possessed by the inquirer is enough to settle a question *Q*, then, when that body of coarse-grained information is a member of *Q*. In other words, the inquirer is in possession of coarse-grained information that constitutes a complete answer to *Q*.

§2.4 That the same point carries over to non-polar questions

We have been exploring a construal of questions whereby they are represented as the sets of their complete answers, which are bits of coarse-grained information (sets of possible worlds).

In the previous section, we have illustrated its workings by looking at polar or yes/no questions. When we want to embed polar questions within the contents of declarative constructions, we use 'whether'-clauses. Examples include the question of *whether the universe started with a Big Bang*, the question of *whether Mary has a girlfriend*, and the question of *whether physicalism is true*.

Such polar questions split the logical space of possibilities in two parts—the *yes* part and the *no* part (one of which might be empty, however). The question of *whether p*, for any *p*, will split the space of possibilities in two parts, namely the part containing the worlds where *p* is the case and the part containing the worlds where *p* is not the case. From there downwards, every subpart of each of the two big parts is also contained in the question—for each of them also counts as a complete answer to the question of *whether p*. Polar questions admit of two maximal answers only (*p* and ¬*p*), and whichever other elements it has are proper subsets of those two maximal elements.[27]

Things get more complex than what this simplistic picture suggests when we try to squeeze questions (construed in this way) out of utterances of natural language interrogatives in general. One important phenomenon here concerns the *presuppositions* that we make when we raise questions in the context we find ourselves in (see §7.2 for the semantic notion of a question's presupposition). Even polar questions often partition a space of possibilities that is smaller than the total set of logical possibilities, in that it is concerned only with possible worlds that satisfy certain

---

[27] As already indicated, they *need* not contain proper subsets of the maximal elements, however, seeing as there are questions such that one of their maximal answers is already the empty set {}. Such answers are contradictions, or some other kind of impossibility. Consider for example the question of whether *Aubrey Plaza (if she exists) is Aubrey Plaza*. One of the maximal answers to that question is that *yes, Aubrey Plaza (if she exists) is Aubrey Plaza*—or *she is identical to herself*. That presumably is the set of all possibilities, in which case the answer *no, it is not the case that Aubrey Plaza (if she exists) is Aubrey Plaza*, amounts to the empty set of possibilities. Where *W* is the set of all possibilities, then, the target question will contain *W* as one of its maximal answers, as well as all of its proper subsets as non-maximal answers. The other maximal answer to it is, however, {}, whose only subset is itself.

presuppositions. A question's presupposition is a necessary condition for *any* of its complete answers.

Consider for example the question of *whether you are still smoking*. The two maximal complete answers to that question are: *yes, you are still smoking* and *no, you are not still smoking* (*you quit it*). In both, the possible worlds where the former is true and the possible worlds where the latter is true, you were a smoker before. But maybe you never smoked to begin with (if you were to say this in response to that question, you would be offering a *corrective* answer, not a complete answer to it). So the question is only partitioning the space of possibilities where you have been a smoker up to now, and those possibilities might not include the actual world, *ergo* it is not the total set of logical possibilities.

On top of that, context will often provide important parameters without which there is no question at all being expressed by the interrogative sentence—although here interrogative sentences are no different from declarative ones, which very often depend on such parameters in order to convey information or to carry content.[28]

'My daughter graduated yesterday', for example, conveys no information if no speaker in particular has produced a token of that sentence, if it wasn't said or written by *someone* (say the sentence was generated by a cosmic accident, or by a series of computers that randomly put words together). The contextual parameters—including the time of the utterance, spatial location, speaker, and possible world where the utterance is made—have to pair up with the *characters* of those expressions to output a question. A character is a function from contexts to contents.[29] For example, the character of the personal pronoun 'I' will output different referents when used by different speakers, in such a way as to account for the difference in content between *my* utterance of 'I love philosophy' and *your* utterance of 'I love philosophy'.

Suppose someone utters 'Who is in favor of it?' at a meeting. In this case, context also has to determine a domain of objects for 'who' to quantify over, and it has to fix the reference of 'it'. Say that 'it' refers here to a new policy, and the people who are under the scope of 'who' are deliberating about it (a time for the present tense in 'is' must also be supplied). Then the target occurrence of 'Who is in favor of it?' expresses a question, and there will be ways of representing it in our framework. For example, say that the people that 'who' is quantifying over are Alice (*a*), Brigit (*b*) and Calum (*c*). Let *pol* be the relevant policy that these people are deliberating about, and let $F(x, y)$

---

[28] The same applies of course to interrogative *complements*—see e.g. Groenendijk and Stokhof (1982, pp. 180-181) on this.

[29] The notion is taken from Kaplan (1989).

be the information that *x is in favor of y*. Then the maximal answers to the question expressed by 'Who is in favor of it?' in this context will be:

(1)	$F(a, pol) \& F(b, pol) \& F(c, pol)$	(they are all in favor of the policy)
(2)	$\neg F(a, pol) \& F(b, pol) \& F(c, pol)$	(all but Alice are in favor)
(3)	$F(a, pol) \& \neg F(b, pol) \& F(c, pol)$	(all but Brigit are in favor)
(4)	$F(a, pol) \& F(b, pol) \& \neg F(c, pol)$	(all but Calum are in favor)
(5)	$F(a, pol) \& \neg F(b, pol) \& \neg F(c, pol)$	(only Alice is in favor)
(6)	$\neg F(a, pol) \& F(b, pol) \& \neg F(c, pol)$	(only Brigit is in favor)
(7)	$\neg F(a, pol) \& \neg F(b, pol) \& F(c, pol)$	(only Calum is in favor)
(8)	$\neg F(a, pol) \& \neg F(b, pol) \& \neg F(c, pol)$	(none of them are in favor)

Each of these is a bit of coarse grained information presented in the notation we have introduced above. For example, the information in line (1) is constituted by the set of possible worlds where Alice, Brigit and Calum are all in favor of the new policy.

So the question expressed by 'Who is in favor of it?' in this case has eight maximal complete answers, as opposed to just two. It splits the space of possibilities into more parts than the polar question expressed by 'Did the universe start with a Big Bang?' does.

And yet we can theorize about this non-polar question the same way we did with polar questions above. For example, we will say that an inquirer's total body of coarse-grained information *i* is enough to settle the target question just in case *i* is a subset of one of (1)–(8). And the inquirer can also have a body of information *i* that constitutes only a partial answer to that question, say, the information that $F(a, pol)$, or that *Alice is in favor of the new policy*, while leaving it open whether Brigit and Calum are in favor of it. In the latter case, *i* isn't a subset of neither of (1)–(8), though it rules out answers (2), (6), (7) and (8).

So the story is more complex, but the pattern is the same. Similar morals apply to a large variety of *when*-, *where*-, *which-,* and *what*-questions (*When is your birthday? Where are my keys? Etc.*), including *mention some*-questions (*What are some places to visit in Rome? Etc.*). Such questions can include indefinitely many maximal answers.

We are not in the business of doing natural language semantics here, however. This framework will be used here rather to build *models of inquiry* and its objects. We don't even have to assume that every question that is represented through our possible-worlds models will find a corresponding natural language expression, much less that *every* natural language interrogative expresses a question in the current sense (we get back to this issue and talk about *pseudo-questions* in §§10.2–10.5).

Even the notion of a *complete answer* is theoretically stipulated here, instead of being grounded on natural language intuitions about which declaratives offer complete answers to the questions expressed by interrogatives. A question in the current sense is taken to be a set containing maximal sets of possible worlds—minimally two—which by definition are complete answers to that very question, plus all of their subsets, which are also complete answers to it (the union of those maximal sets must equal the total set of possible worlds that is partitioned by the question's maximal answers, namely, the set of possible worlds where all the question's presuppositions are true).

What we are doing now is we are exploring the explication of what it is for a subject to have information that is enough to settle a question introduced above. The next chapter labors the point further and finds limitations with the model of inquiry it engenders. Those limitations point, however, to challenges that stand in the way of a better understanding of armchair inquiry itself.

Chapter 3

§3.1 Where a model of inquiry is fleshed out

Our theoretical constructs of what a question is and of what it takes for an inquirer to have enough information to settle a question provide us with a *model* of inquiry.

Our most basic theoretical constructs are the so-called *possible worlds*, or ways the world might be—if we want, actual and counterfactual scenarios.[30] These fictional or abstract objects are then used to construct bodies of coarse-grained information, which are simply sets of possible worlds. The coarse-grained information that *p* is again the set of worlds where '*p*' is the case (where '*p*' has the meaning that it has in *our* mouths, in the context that *we* utter that sentence).

Such sets allow us to capture the coarse-grained information that is possessed by an inquirer, on the one hand, and what the answers to a question are, on the other. Since a question is here represented as the set of its complete answers/the downward closed set of its maximal answers, questions themselves are built out of such bits of coarse-grained information. All in all, then, the model allows us to represent situations where an inquirer already has information that is enough to settle a given question, distinguishing them from those situations where she doesn't. In the former kind of situation, so the model goes, the inquirer's coarse-grained information is a member of the question, and in the latter one it isn't.

Successful information-gathering inquiry is neatly accommodated here as the elimination of previously uneliminated possibilities. That is, when the inquirer acquires new information, she transitions into a new informational state where the space of possibilities that constitutes her total body of information is a *proper subset* of the space of possibilities that constituted her total body of information before the acquisition of that data.[31]

According to this picture, information-gathering inquiry is successfully concluded when, after eliminating certain possibilities, the inquirer is left with a body of

---

[30] Such theoretical constructs are deployed in philosophy and semantics for a number of different purposes. See Menzel (2021, §2) for the many different takes one might have regarding the ontology of possible worlds.

[31] This picture of inquiry is famously put forward by Stalnaker (1984). He also tries to accommodate what we are calling 'armchair inquiry' here with a metalinguistic strategy, where the information we learn via deduction is concerned with the properties/relations of linguistic expressions. When it comes to armchair inquiry, the proposal developed in this chapter and the next is an alternative to Stalnaker's.

coarse-grained information that is one of the answers to the question she was inquiring into—that is, a member of the question itself. Information-gathering inquiry is an activity designed to eliminate possible ways the world could be to the inquirer, until no more elimination is needed.

Suppose an art collector named Jane Finder finds a previously unknown surrealistic painting of Frida Kahlo. The painting is undated, however, and now Finder wants to find out *when (in which year) Frida Kahlo painted it*. Suppose at time t_1 the information available to Finder allows her to narrow the possible years down to those in the interval 1932–1934. At this point, Finder's total body of coarse-grained information can be divided between (a) possible worlds where the painting was painted in 1932, (b) possible worlds where it was painted in 1933, and (c) possible worlds where it was painted in 1934. So if Finder were to learn going forward that *Frida Kahlo painted that painting in February 1933*, say, she would thereby have a total body of coarse-grained information that is enough to settle her question, since that total body of information would then constitute a complete answer to it.

But suppose instead that she learns first, at t_2, that *Frida Kahlo wasn't able to paint any paintings in 1934* (this is fictional, not based on facts). Then what happens is that Finder's total body of information at t_2 eliminates possibilities that were left open at t_1, namely, all those possible worlds mentioned in (c) above, though she still doesn't have a complete answer to her question at t_2. She has made *some* progress with her process of inquiry at t_2, in that she can now rule out a number of complete answers to her question that she couldn't rule out before. But there is still more progress for her to make here—she needs to gather more information.

Such is the way in which information-gathering inquiry goes forward. It will reach its successful end when, at some time t_n, Finder ends up with a body of coarse-grained information that is a subset of (a) or a subset of (b). At that point, the question will be *closed* relative to the information possessed by her, and there will no more progress for her to make in that regard. As long as her coarse-grained information has both, worlds belonging to (a) and worlds belonging to (b), however, she doesn't yet have information that is enough to settle her question, or her question is still left open relative to the information possessed by her.

The progress of information-gathering inquiry is tracked by updates in the coarse-grained information possessed by the inquirer at different times. The phenomenon can be captured through a combination of the canonical frames from epistemic logic, each containing a total set of possible worlds and an accessibility

relation, and an update operation on those frames, as in dynamic epistemic logic.[32] The accessibility relation will tell us which possible worlds are left open or are compatible with the inquirer's information at any particular time. When we update a frame on a new piece of information *p*, that accessibility relation will be updated accordingly, in that possible worlds where *p* is false are not accessible relative to the new accessibility relation (even though they might have been accessible relative to the old one).

As the inquirer learns new things, the space of possible worlds that are left open by the information possessed by her gets chopped down further, for less and less worlds are accessible to her as time goes by. If we represent that space of possibilities at a given time as a circle within a larger space of possibilities, that circle gets smaller and smaller as the inquirer acquires more coarse-grained information.

Paired with some auxiliary assumptions, the model vindicates instrumental norms of inquiry that sound quite plausible. Assume, for example, that when the inquirer knows that *p* at *t*, then at *t* the inquirer counts as having the coarse-grained information that *p*—so every possible world that is compatible with the inquirer's total information at *t* is a world where *p* is the case (in the model, her accessibility relation only relates worlds where *p* is true to the world she is in at *t*). Given that much, we can now state:

(W) If one already knows *Q*, then it would be a waste of one's time and resources to collect more information in an attempt to settle *Q*.

Now where the 'should' is the 'should' of *instrumental rationality*, we can use (W) to ground the following instrumental norm of inquiry:

(N) One shouldn't collect more information in an attempt to settle *Q* if one already knows *Q*.

(This norm follows from a more general norm—the Ignorance Norm of Inquiry—which is discussed in §6.3 below. In contrast to (N), the more general norm is concerned not only with the collection of new information, but also armchair inquiry. It is called that way because it says that one shouldn't inquire into a question *Q* unless one is *ignorant* of which of *Q*'s maximal complete answers is true).

Here again, '*Q*' is used as a placeholder for interrogative complements such as 'whether the water is hot', 'where we came from', 'who declared independence', etc.

---

[32] See Pacuit (2013a) and (2013b) in this sequence for the gradual building of this theoretical package, and also van Ditmarsch, van der Hoek and Kooi (2008).

You shouldn't collect more information in an attempt to settle the question of *whether you like bananas*, say, if you already know *whether you like bananas*. And one shouldn't collect more information in an attempt to settle the question of *who let the dogs out* if one already knows *who let the dogs out*. Etc. How does the possible-worlds model of inquiry justify (N)?

As we saw, what information-gathering inquiry does in this picture is to chop the space of possibilities that are open to the inquirer down to smaller and smaller subsets. Eventually, the inquirer eliminates as many possibilities as needed to end up with a body of information that is enough to settle a given question—in which case any further elimination of possibilities will be *otiose*, as far as the goal of settling that question is concerned. For, when the inquirer reaches that stage, her body of coarse-grained information (a set of possible worlds) is *already* one of the question's answers. Any proper subset of it will continue to be just that: one of the question's answers. Since the situation cannot get better than that with respect to that question in particular, there is no more point in collecting more information that bears on it anymore—the question is already closed relative to the inquirer's total body of information. Finally, when the inquirer comes to know one of the question's complete answers, she reaches exactly that stage with respect to that question: she has information that is enough to settle it.

This vindication of an apparently true norm of inquiry notwithstanding, we now have to look at some of the important limitations of this model of inquiry.

§3.2 That the possible worlds model does not capture crucial aspects of armchair inquiry

In the possible worlds model of inquiry, questions are construed as sets of coarse-grained information—the set of *complete answers* to that question—and the information possessed by the inquirer is also construed as a body of coarse-grained information, which is some set of possible worlds.

We saw that this model allow us to represent situations where an inquirer already has information that is enough to settle a given question. A distinction is thereby drawn between the questions that the inquirer already has enough information to settle and the questions that she does not already have enough information to settle.

But the model does *not* distinguish the questions that the inquirer is able settle through armchair inquiry from the questions that the inquirer is not able to settle through armchair inquiry, even though she already has information that is enough to settle them. Previously (§1.4) we saw that, for any inquirer at any particular time, any question *Q* will satisfy exactly one of the following descriptions:

(I^+S^+)	The inquirer has enough information to settle *Q*, and also the skills that are needed to settle *Q* on the basis of that information.
(I^+S^-)	The inquirer has enough information to settle *Q*, but not the skills that are nee to settle *Q* on the basis of that information.
(I^-)	The inquirer doesn't have enough information to settle *Q*.

Now the questions that the inquirer is able to settle through armchair inquiry, as we saw, are questions of type (I^+S^+). And the questions that the inquirer will be able to settle only after she acquires new information are always questions of type (I^-). Sandwiched between them, there are also questions of type (I^+S^-), which are *neither* questions that the inquirer is able to settle through armchair inquiry, *nor* questions that the inquirer will be able to settle only if she acquires new information. Questions of type (I^+S^-) are not questions that the inquirer is already in a position to settle—but that is because she lacks the skills that are needed to settle them on the basis of the information available to her, rather than because she lacks information that is enough to settle them.

The possible worlds model of inquiry, however, lumps questions of type (I^+S^+) and question of type (I^+S^-) together (for any inquirer at any time) within a single category. It distinguishes questions of type (I^-) from the rest without distinguishing the rest itself. The presence or absence of cognitive skills that are needed to settle a question is not captured by that model. The only feature that it captures is the inquirer's having or not having information that is enough to settle a question. And that means that the model doesn't distinguish between the questions that the inquirer is already in a position to settle from those that the inquirer is not already in a position to settle.

For the questions that the inquirer is already in a position to settle are the questions of type (I^+S^+), and the questions that the inquirer is not already in a position to settle are the questions that are *either* of type (I^+S^-) *or* type (I^-). As far as the powers of discrimination of the model goes, however, there is just one big blob of questions including both questions of type (I^+S^+) and questions of type (I^+S^-), on the one hand, and another one including questions of type (I^-), on the other.

Here is another angle from which we can point to the limitations of this model. It doesn't allow us to capture some essential aspects of armchair inquiry, however good a model of information-gathering inquiry it turns out to be. Let us labor the point in more detail.

Consider the following kind of situation where armchair inquiry comes to a successful end: the inquirer settles her question by competently *deducing* an answer to it from the information possessed by her (for the sake of illustration, think of deduction

that tracks *logical* entailment relations here). The problem is that the possible-worlds model of the inquirer's cognitive situation with respect to that question *before* and *after* she has settled her question in this way may remain exactly the same—for her coarse-grained information need *not* have changed in the transition between from the initial to the final stage of her process of inquiry.

In fact, to the extent that the inquirer was able to settle her question through sheer armchair inquiry in the first place, without acquiring any new coarse-grained information, she must have already had an answer to her question (an element or member of that question) *before* making the deduction. Something has changed in the cognitive life of the inquirer, however, after she has deduced the answer to her question. That change isn't tracked by the possible-worlds model of her situation. It is as if the inquirer's successfully concluding her armchair inquiry in this way makes no difference to her cognitive life.

The issue is most clearly illustrated through cases where deduction gives the inquirer something of an 'aha moment'. The conclusion really strikes the inquirer as something *new*, or something that she has learned only after thinking through the problem (and not before that). For example, suppose you have the following bits of information about spies *A*, *B* and *C* at time *t*:

(a) *Spy A is using binoculars to spy on spy B*.
(b) *Spy B is spying on spy C.*
(c) *Spy C is spying on spy A, but C is not using binoculars.*

Now you are prompted to answer the question: *Is there a spy who is using binoculars to spy on a spy who is not using binoculars*? You start inquiring into that question using the information you have. There will be more information possessed by you other than (a)–(c), of course, such as that *A, B and C are all in the same block*, etc. But the information in (a)–(c) is already *enough* to settle your question. In order to get to the right answer, you may reason as follows:

Either B is using binoculars to spy on C or not.
If B is using binoculars to spy on C, and C is not using binoculars, then there is a spy who is using binoculars to spy on a spy who is not using binoculars.
And if B is not using binoculars to spy on C then, since A is using binoculars and A is spying on B, then again there is a spy who is using binoculars to spy on a spy who is not using binoculars.
So, either way, there is a spy who is using binoculars to spy on a spy who is not using binoculars.

So there is a spy who is using binoculars to spy on a spy who is not using binoculars—case closed.

Think of each line from top to bottom above as recording the contents of your thoughts up to the time where you settle your question at the end. It seems that by the time you reach the last line you learn something new—*now* you have settled your question. You were not there yet at the time corresponding to the first line, at the beginning of your thinking process. You had not yet settled your question then.

But now consider the coarse-grained information that you had at the start of your reasoning process and compare it to the coarse-grained information you had at the end of it. In the beginning, your coarse-grained information was a set of possible-worlds where all of (a), (b) and (c) were true. At the end of your process of (armchair) inquiry, your coarse-grained information doesn't rule out any of (a), (b) or (c). Quite to the contrary, it better *still* be a body of coarse-grained information where all of (a), (b) and (c) are all true, so as to guarantee that the conclusion you arrived at is true.

Your conclusion doesn't chop the space of possibilities open to you down to a smaller set—it was true at all the possible worlds open to you all along. You *had an answer to your question* all along, your course-grained information has not changed after the inference. So the possible-worlds model of your situation before and after your 'aha moment' is exactly the same—no change to be seen *there*.

§3.3 That armchair inquiry is however a puzzling phenomenon in and of itself

We have just pointed to some important limitations of the possible-worlds model of inquiry. When we construe questions as sets of answers, and we construe *both* the answers to a question and the information possessed by an inquirer as sets of possible worlds, we do get to capture distinguishing features of information-gathering inquiry (progress in inquiry via the acquisition of information that eliminates previously uneliminated possibilities), but we miss out on important features of armchair inquiry (progress in inquiry via deductive reasoning).

That model of inquiry lumps together questions that we would otherwise want to keep separate—namely, questions that the inquirer is already in a position to settle and questions that he is not already in a position to settle, even though he has information that is enough to settle them. The inquirer's epistemic standing with respect to questions of the former kind is importantly different from his epistemic standing with respect to questions of the latter kind.

Furthermore, no update in the set of possible worlds that is compatible with the inquirer's information is able to track the changes that occur in his cognitive life after he settles a question by deducing an answer to it from the information available to him. It

is as if deduction made no difference to inquiry. We are going to explore the limitations of the possible-worlds model of inquiry not only from the angle of the problem of the informativeness of deduction, but also from the angle of the taxonomy of questions described above. We start with the former, and come back to the latter in §4.3.

It would be unfair to strictly attach the problem of the informational yield of deduction to the possible-worlds model of inquiry. For the problem can be motivated without resource to the theoretical constructs of that model.[33]

Armchair inquiry is at its best when the inquirer deduces answers that sound *new* to her in some sense. But, in another sense, those answers must not be new at all, on pain of the inquirer's deduction being invalid: if the conclusion *says more* than the conjunction of the premises (the information possessed by the inquirer), then there are situations where all the premises are true and the conclusion is false.[34] To 'say more' is in effect to be more informative. For example, *Willard was a philosopher and a spy* says more than *Willard was a philosopher*, and *Ruth was a female logician* says more than *Some logicians are female*.

In order for our armchair inquirer to even be in a position to settle her question via competent deduction, then, she must in a sense already have the information that constitutes the answer to her question. Otherwise, she would need to gather more information first in order to settle her question—empirical research would be called for and armchair inquiry wouldn't be enough.

So, when the inquirer settles her question via deduction, her answer *better* not be new information, relative to the information she already had before the deduction. Otherwise, the deduction would not be competent—its conclusion wouldn't follow from its premises. We can also give the point a formal gloss. Where q follows from premises $p_1,\ldots,p_n$, the conjunction (p_1 &…& p_n) is *equivalent* to (p_1 &…& p_n & q). Adding the conclusion q to (p_1 &…& p_n) only brings sand to the beach. Throughout the history of philosophy, philosophers and logicians have also used the containment metaphor: in

---

[33] Hintikka (1973) observes that the classical semantic information theory put forward by Bar-Hillel and Carnap (1952)—which is translatable into the theory of coarse-grained information adopted here—doesn't allow us to capture the information-gain that is sometimes provided by deduction. Hintikka, however, does not look at the problem as if it was confined to the classical semantic information theory, and he acknowledges a more general form of the problem which he calls the 'scandal of deduction': if logical truths do not say anything informative, and deductive inferences correspond to logical truths (via the deduction theorem), then how come can deductive inference yield new information? (1973, p. 222). See also Sequoiah-Grayson (2007) for criticism and discussion of the way Hintikka went about dealing with the problem.

[34] And so, Wittgenstein in the *Tractatus*, 5.14: 'Where a proposition follows from another, the latter says more than the former, the former less than the latter' ('Folgt ein Satz aus einem anderen, so sagt dieser mehr als jener, jener weniger als dieser').

order that the conclusion follows from the premises, the former must already 'be contained' in the latter.[35]

The problem that presents itself is that of explaining in what sense a deduced answer can be *new* information to the inquirer (so that the inquirer may *learn* new things through deduction) and in what other sense it is *not new* information, with respect to the information already possessed by the inquirer (so that the deduced answer indeed *follows* from that body of information).[36] We have some work to do here, independently of our theorizing about inquiry with the possible-worlds model.

Before proceeding, however, it will be helpful to us—and it will help us prepare for the next steps of this investigation—to say something about what the problem with the possible worlds model of inquiry *is not.*

In this model, questions are the sets of their complete answers, both the answers to a question and the information possessed by the inquirer are bits of coarse-grained information, and bits of coarse-grained information are sets of possible worlds. The learning that takes place in empirical search is modeled by updates in the frames that represent the inquirer's situation, where previously uneliminated possibilities within the inquirer's body of coarse-grained information come to be ruled out.

But the inquirer's validly deducing an answer to her question, we saw, does not change the coarse-grained information possessed by her. That captures the sense in which deductive armchair inquiry does *not* give the inquirer new information, so that the deduction of the answer on the basis of the information possessed by the inquirer is indeed valid. But it misses out on the sense in which deductive armchair inquiry *does* give the inquirer new information, so that she learns something new and finally settles the question that she was inquiring into at first.

The problem with our possible-worlds model of inquiry is *not* the so-called problem of 'logical omniscience', which is often treated as a bug of canonical systems of epistemic and doxastic logic.[37] For our models do *not* say that the inquirer believes or knows everything that logically follows from what she believes or knows, including all logical truths. In fact, our model is silent about belief and knowledge, since in its intended interpretation it merely represents the coarse-grained information that is possessed by the inquirer, and how the status of a question (open or closed) stands

---

[35] See for example De Morgan (1847, p. 254), who adds later, however, that 'the presence of the premises in the mind is not necessarily the presence of the conclusion' (*idem*). If we were to use De Morgan's terminology, we would describe the task that is presented in the next paragraph as that of explaining *in which sense* is the conclusion already 'contained' in the premises, and in what sense can the premises 'be present in the mind' without the conclusion 'being present in the mind'.

[36] On this point, see also Dummett (1978) and, more recently, Jago (2013).

[37] See Hintikka (1962, pp. 29–31), Stalnaker (1991), and also Yap (2014) for discussion.

with respect to that information. It does not assume that if the inquirer possesses the coarse-grained information that *p* then she believes/knows that *p*.

The application conditions for 'possesses the coarse-grained information that…' must be *looser* than those for 'knows that…', for example. To illustrate, say that we use knowledge ascriptions to determine the coarse-grained information that is possessed by an inquirer—a method that goes as follows: (i) The coarse-grained information possessed by an inquirer is a set of possible worlds *i* such that *w* belongs to *i* if and only if, for *every p* such that the inquirer knows that *p*, *p* is true in *w*; (ii) We say that the inquirer possesses the coarse-grained information that *p* if and only if *p* is true at all the members of such a set *i*. The inquirer's set of uneliminated possibilities is in this way determined by what she knows. Now it follows from this method that if the inquirer has the coarse-grained information that *p,* and *p* entails that *q*, then she also has the coarse-grained information that *q* (more below). But of course this is perfectly compatible with the claim that the inquirer can know that *p* and not know that *q*, even though *p* entails that *q*. For (i) and (ii) do not entail that the inquirer knows that *p* whenever she has the coarse-grained information that *p*.

In §2.3 we talked about fine-grained information, and we said that the coarse-grained information possessed by an inquirer is the set of possible worlds where the fine-grained information possessed by her is true. In the previous paragraph, we thought of fine-grained information as *that which the subject knows*, which was one of the suggestions from §2.3. But this was just an example of a way of refining the general idea. The general thrust of it is: the fact that coarse-grained information is closed under entailment does not entail that fine-grained information is closed under entailment. Accordingly, our models of inquiry do not depict the inquirer's *fine-grained* information as being closed under entailment—in fact, they do not contain any constructs that stand for that kind of information to begin with.

Fine-grained information, we also saw in that section, is intimately connected with *how the subject represents* the world as being, or what her *representational states* are. The inquirer has or possesses a certain body of coarse-grained information in virtue of the fine-grained information that is possessed by her in virtue of the fact that some of her representational states carry that coarse-grained information. Her representational states feature *vehicles* of information such as pictorial, iconic and

symbolic representations.[38] Such vehicles constitute the fine-grained information themselves. To every body of fine-grained information (say, what the inquirer knows/ knows explicitly), there corresponds a body of coarse-grained information (say, the set of worlds where every proposition that is known/explicitly known by the inquirer is true). Many ascriptions of coarse-grained information ascribe information that is only *implicit* in the subject's representation of the world.[39]

Accordingly, it seems that it is new fine-grained information that the deductive inquirer acquires after she settles her question via deduction. The coarse-grained information remains the same old one, so that she can safely add new fine-grained information—new representational vehicles—to her total representation of how the world is like. That is a more or less natural way of addressing the problem of deductive armchair inquiry that we saw above. What that kind of inquiry does is make explicit what was only implicit before. The next section expands upon this theme.

§3.4 That competent deduction affords new *fine-grained* information while preserving coarse-grained information

The problem that we raised in the previous section was that of explaining in what sense a deduced answer can be *new* information to an inquirer, and in what sense it is *not new* information with respect to that inquirer's already possessed information.

Hints to a solution occurred to us by the end of the previous section: even though the deductive inquirer's coarse-grained information is not new, her fine-grained information is new. We were thinking of fine-grained information itself as being constituted by representational vehicles of a certain sort.

That allows coarse-grained information to be closed under entailment when fine-grained information isn't. If the inquirer possesses fine-grained information that p, then she possesses the coarse-grained information that p. But coarse-grained information extrapolates fine-grained information. It is *not* the case that if the inquirer

---

[38] What we say here is supposed to be compatible with deflationist accounts of mental representation such as the one put forward by Egan (2020), who gives an account 'of what it is to function as a representation in an explanatory account of a cognitive capacity' (p. 40). Furthermore, at a physiological level of description, the types of representation appealed to might all boil down to a certain type of neural representation or patterns of neural activation. See Hutto and Myin (2014), Bechtel (2016), Thomson and Piccinini (2018) for different takes on the status of neural representations. We won't take a stand on these issues.

[39] And so van Benthem and Velázquez-Quesada (2010), who try to set apart implicit and explicit knowledge within a system of dynamic epistemic logic, point out that in the distribution axiom of normal epistemic logics, $K(p \supset q) \supset (Kp \supset Kq)$ ($\supset$ is the material conditional), 'the K operator really just describes implicit semantic information of the agent' (2010, p. 246). They remark subsequently that this implicit semantic information *definitely* has the relevant closure property (quite generally, it is closed under entailment). But the same won't apply to what the agent knows in an explicit manner.

possesses the coarse-grained information that *p*, then she possesses fine-grained information that *p*. She can possess the coarse-grained information that *p* without possessing fine-grained information that *p*. The role of deduction would then be to close this gap in the subject's representation of how the world is like, allowing her to better pair her fine-grained information to her coarse-grained information. What is new to the inquirer after valid deductive inferences is new representational vehicles.

Armchair inquiry now becomes less puzzling. Suppose the inquirer is about to settle her question via valid deduction from the information possessed by her. In terms of what we saw in the previous chapter, she has the cognitive skills that she needs in order to reliably form a belief in one of the question's answers on the basis of the information possessed by her. But at this point she hasn't done that yet. Since her deduction will be valid, the question she is inquiring into is already effectively *closed* relative to the coarse-grained information possessed by her, even before the deductive inference takes place. Her coarse-grained information is in effect one of the possible answers to her question—though of course she herself doesn't know this (she is like the slave boy in Plato's *Meno*, who didn't know he had the answer to the question of *how one can double the area of a square* before inquiring into that question under the guidance of Socrates' speech acts).

If our inquirer's question were not already closed relative to her coarse-grained information, after all, she would not be in a position to settle her question through a valid deductive inference from the information possessed by her to begin with. Now she makes the deduction and thereby forms a belief in one of the question's answer. It is at this point that her *fine-grained* information changes, so that she learns something new. She must have done some kind of transformation in the vehicles of her thought. We might want to put it like this: the inquirer has found a new *guise* with which to present coarse-grained information she had all along.

Deduction puts old things under new guises. The new guises in turn create yet further inferential and decision-making opportunities to the inquirer, for it gives her new vehicles with which to cognize and prompt behavior. In this way, fine-grained information (the *new* information) can keep growing and enriching the inquirer's representation of reality, while her coarse-grained information (the information that is *not new*) remains steady, so as to ensure that the inferential steps whereby she adds new fine-grained information are valid, as deductive inferences should be.

This story also squares well with things that epistemologists have written about the expansion of knowledge via (competent, reliable) deduction. Much of the literature has concentrated on principles of 'closure' of knowledge under entailment/known

entailment.[40] To the extent that some such principle of closure is accepted, it has the reasoner *competently deduce* the conclusion from known premises. Otherwise, how could knowledge of the premises confer knowledge upon the conclusion? John Hawthorne, for example, writes: 'If at *t*, I know that *p* and I know that *p* entails *q*, I may still have to *do* something—namely, perform a deductive inference—in order to come to know that *q*. Until I perform the inference, I do not know that *q*' (2004, p. 32).

Whether or not the reasoner performs the deductive inference, the coarse-grained information that comes with her knowledge—the set of possible worlds where the propositions that are known by her are all true—remains the same. What performing the deductive inference does is give the reasoner new fine-grained information. In coming to know that *q* on the basis of deduction from known premises he acquires fine-grained information that *q*.[41]

That the deductive inference was performed *competently* means, among other things, that the reasoner already had the coarse-grained information that *q* even *before* he drew his conclusion (again, having the coarse-grained information that *q* doesn't yet entail knowing that *q*). That is why he still had to *do* something in order to come to know that *q*, to make explicit what was only implicit. He had to *process* his fine-grained information in a certain way, which here means processing representational vehicles (information processing as vehicle processing).

The processing of fine-grained information that competent deduction consists of provides our reasoner with new representations, which have a new format, while preserving his coarse-grained information. The deductive inference is competent (*qua* deduction) just in case the conclusion that *q* is indeed true in all the possible worlds that are compatible with the reasoner's information even before he performs the inference. When the reasoner adds the new fine-grained information that *q* to her representation of reality, no possible world that was part of his coarse-grained information before the inference is ruled out. And only *then* does he come to know that *q.* Before all this, he was only in a position to know it. The processing of fine-grained information is what was missing.

What kinds of *representational vehicles* should we be thinking of here, more exactly? We have to leave the issue somewhat open, seeing as the armchair methods available to us here are not the most appropriate ones to settle it.

Luckily, however, the topic has been covered by cognitive scientists. Those working within the 'mental logics' school in psychology of reasoning posit symbolic

---

[40] See Luper (2020) for an overview of this literature.

[41] We don't need to assume here that the inquirer has the fine-grained information that *p* if and only if she knows that *p*. We just assume that acquiring new knowledge that *q* via deduction entails acquiring the fine-grained information that *q*.

representations or mental sentences, language-of-thought style, and they argue that reasoning consists of performing operations over representations of that kind.[42] 'Mental models' theorists, on the other hand, postulate iconic representations, whose spatial properties and relations are mapped onto the properties and relations of the things they are about (representation via similarity).[43] We reason by building mental models of situations and checking what else is true of the situations we have mentally modeled. Both schools of thought appeal to their respective kinds of representation in order to explain *how* we reason, as well as to account for the systematic mistakes that we make in reasoning, our cognitive limitations in performing certain reasoning tasks (for example, regarding how much information can be stored in our working memory).

Unsurprisingly, others will maintain that, at the level of cognitive processing, reasoning can consist of the manipulation of *both* kinds of representations, symbolic and iconic ones.[44] And there should still be other ways of describing the process of reasoning as one that manipulates representational vehicles, say, understood as patterns of neural activity.[45] But we don't have to decide between the options here, or see if/how they can be integrated into a single picture. The point is that it doesn't seem to be empirically inadequate to talk of reasoning as manipulation of *some* kind of representational vehicle, which is all we need. The disagreements among psychologists working on the topic are not about *that* shared assumption.

We now need to look at some of the structural features that such representational vehicles must have in order to allow fine-grained information to play the role it is supposed to play here (to be the *new* information that is gathered through

---

[42] See Rips (1994) and also Braine & O'Brien (1998). Here, the faculty of reasoning is conceived as something like an automated theorem-prover, only it 'proves' not only standalone theorems, but also things that follow from premises. Our reasoning system extracts the logical forms of premises, which are symbolic representations, and applies rules to those premises—rules whose antecedent conditions describe exactly those forms—to derive further conclusions.

[43] See Johnson-Laird and Byrne (1995), Johnson-Laird and Khemlani (2013). Where reasoning through *sentences* is concerned, mental model theorists explain it roughly as follows: (a) we build mental models for the linguistically formulated premises, where these models stand for the possibilities where those premises are true, (b) we read further linguistically formulated conclusions off those mental models, sometimes conclusions that hold in all the mental models of the premises, other times conclusions that only hold in a large proportion of them, and (c) in more demanding situations, we also search for counterexamples, or for models of the premises that are not models of the conclusion. This is again their description of inference when premises and conclusions are presented in an explicit linguistic format to the subject; at a more fundamental level, reasoning is the process of transforming or updating one's mental models.

[44] Relatedly, Goel (2005) marshals neurological evidence to the effect that, whereas some reasoning tasks engage mostly parts of the brain that are responsible for the processing of spatial representations, other reasoning tasks engage mostly parts of the brain that are responsible for language-processing.

[45] On the topic of how liberal we can be regarding which neural processes count as representational, see Martínez and Artiga (forthcoming).

deductive armchair inquiry). This will also give us pointers as to how an alternative model of inquiry should look like if it were to capture the contribution that deduction makes to inquiry. We then revisit the problem of splitting questions into three categories, for every inquirer: the ones that she is already in a position to settle, the ones that she has enough information to settle but lacks the cognitive skills to do so, and the ones that she doesn't yet have enough information to settle.

Chapter 4

§4.1 Where the distinction between fine-grained and coarse-grained information is further explicated

In the previous chapter, we saw that the possible worlds model of inquiry is not able to capture the contribution that deduction makes to armchair inquiry. The problem gained a more model-independent face through the observation that deduction must *not* give the inquirer new information if it is to be valid, in one sense of 'information', whereas it sometimes provides the inquirer with new information, in another sense of 'information'.

Our first step towards a solution was to use the notion of *coarse-grained* information to play the former role, and the notion of *fine-grained* information to play the latter one. Fine-grained information, in contrast to coarse-grained information, is constituted by *representational vehicles*.

A piece of fine-grained information is a pair of a vehicle and the coarse-grained information that is carried by that vehicle. What deduction does is add such vehicles to the inquirer's representation of how the world is like, her internal model of reality. The role of deduction is to enrich the inquirer's stock of fine-grained information. When the deduction is valid, it does that without changing the inquirer's coarse-grained information, to the inquirer's own benefit (no risk taken by her here).

In §3.4, we have left it somewhat open exactly what type of representational vehicle must a piece of fine-grained information be constituted by. We don't necessarily take them to be *sentences* that are tokened in the inquirer's thoughts (mental reproductions of sounds or inscriptions), for example. We just need it to be *some* kind of vehicle, not necessarily any particular kind of vehicle.[46]

Such representational vehicles, however, must satisfy certain structural conditions in order to constitute what we are calling 'fine-grained information'. These conditions are largely concerned with the relationship between coarse-grained and fine-grained information. The issue is a delicate one, partly due to the fact that we embed sentences (which are themselves representational vehicles) in our *ascriptions*

---

[46] It should be possible for us to do away with our apparent reification of representational vehicles as discrete parts of one's total representation of reality—say, by paraphrasing the sentences that seem to commit us to the presence of such vehicles in our domain of objects with a language that talks about the phenomenon of represent*ing* the world as being a certain way, or a language whose expressions stand for physical *events* of representation, rather than to objects that represent. The paraphrase mechanism should allow us to maintain the structural conditions that will occupy us next.

of information-possession to inquirers. It is easy here to create the impression that the vehicle that *we* use in our ascription is supposed to be the same as the vehicle that *the inquirer* thinks with.

The first thing to note here is that we have indiscriminately used the same complement 'that *p*' to ascribe both, coarse- and fine-grained information possession. This linguistic habit is a convenient one, if only for its conciseness. Attention must be paid, however, to how the truth-conditions of the two kinds of ascriptions relate to each other and how they can come apart. They do come apart (when they do) because the semantic contributions of 'the coarse-grained information that *p*' and 'the fine-grained information that *p*' respectively are not the same, despite the fact that they embed the same 'that'-clause.

We have assumed that the acquisition of fine-grained information moulds the inquirer's coarse-grained information, in that her coarse-grained information is the set of possible worlds where her fine-grained information is true/accurate. We can say that an inquirer's fine-grained information *determines* her coarse-grained information, in that the representational vehicle that constitutes her fine-grained information that *p carries* the coarse-grained information that *p.* Coarse-grained information gets imported into the inquirer's mind by means of her fine-grained information, so to speak.[47]

So the truth-conditions for 'the inquirer has the fine-grained information that *p*' must entail the truth-conditions for 'the inquirer has the coarse-grained information that *p*'. But we have assumed that inquirers can be in possession of a bit of coarse-grained information without having a corresponding bit of fine-grained information. So the truth-conditions for 'the inquirer has the coarse-grained information that *p*' must *not* entail the truth-conditions for 'the inquirer has the fine-grained information that *p*'. After all, what the deductive inquirer is supposed to learn is new fine-grained information, while keeping her coarse-grained information the same. This strategy for explaining what the value of deduction is to armchair inquiry would be useless if having the coarse-grained information that *p* already entailed having the fine-grained information that *p*.

To flesh this out in a more precise manner, let i_w be the set of possible worlds that constitutes the inquirer's body of coarse-grained information at possible world *w*.[48] So 'the inquirer has the coarse-grained information that *p*' (as uttered at context *c*) is true at *w* if and only if '*p*' (as uttered at that same context *c*) is true in all the members

---

[47] Similarly, Cresswell (1975, p. 25) suggests saying that a proposition *determines* a set of possible worlds, without being identical to them.

[48] We should really be thinking here of the coarse-grained information that is possessed by the inquirer at a certain world *at a certain time*. But we omit reference to time in the interest of avoiding clutter—so let us think of *w* as being centered on a particular time.

of i_w (for the sake of simplicity, we leave the pairing of the contextual parameter implicit from now on).[49]

Those are then the truth-conditions for ascriptions of coarse-grained information. They give our ascriptions of coarse-grained information the following 'closure' property, again: if the inquirer has the coarse-grained information that *p*, and *p* entails that *q*, then she also has the coarse-grained information that *q*. That the inquirer 'has' the coarse-grained information that *p* just means that her representation of how the world is like carries the coarse-grained information that *p*—and it thereby carries the information that *q*, for any *q* that is entailed by *p*.

We can also formulate our truth-conditions using *valuation functions* from the canonical models of modal logic (plus set theory). Let *v* be such a valuation function. It maps sentences of our language into sets of possible worlds—so that *v*('*p*') is the set of worlds where '*p*' is true. Now we can just say that 'the inquirer has the coarse-grained information that *p*' is true at *w* just in case $i_w \subseteq v(\text{'}p\text{'})$. That is: just in case the set of possible worlds that constitutes the inquirer's body of coarse-grained information at *w* is a subset of the set of worlds where '*p*' is true.

That '*p*' is true at all the members of i_w is also a *necessary* condition for the truth of 'the inquirer has the fine-grained information that *p*' at *w*. This secures the entailment relation from 'she has the fine-grained information that *p*' to 'she has the coarse-grained information that *p*'. But it cannot be a *sufficient* condition for that, so that the entailment relation in the other direction doesn't hold.

What else is needed for the possession of the fine-grained information that *p*? Following suit on our observations from above, it must involve the incorporation of a certain kind of vehicle into the inquirer's representation of reality. That vehicle must of

---

[49] By 'context' we mean again some tuple containing a speaker, time, possible world, and possibly other elements still. To see why the qualification matters, consider my utterance of 'You have the coarse-grained information that *I wrote this book*'. What I am saying here is not that your coarse-grained information is constituted only by worlds where the sentence 'I wrote this book' *as uttered by you* is true.

course carry the coarse-grained information that *p*, so that it represents the world as being such that *p*.[50] Could that vehicle simply be the declarative sentence '*p*'?

It could be, of course, but it need not be. For, even if representational vehicles are sentences—or at least *some* of those vehicles are sentences—maybe our sentence '*p*' is not formulated in a language that our inquirer can think with (the one we are ascribing information possession to). Maybe our inquirer doesn't even master the language with which '*p*' is built, in which case '*p*' is not the representational vehicle through which she represents the world as being such that *p*.

For example, we might want to describe the ancient geometer Euclid as deducing that *the sum of the angles of a triangle is always equal to 180 degrees* from his axioms. According to the theory we started to develop above, that means that Euclid has added *some* vehicle to his representation of reality that carries the coarse-grained information that *the sum of the angles of a triangle is always equal to 180 degrees*. But Euclid didn't speak English. So the vehicle of his thought was not (a token of) the sentence 'The sum of the angles of a triangle is always equal to 180 degrees'. What should the vehicle that constituted Euclid's fine-grained information be, then, or how should we think of it?

Our problem here is of course reminiscent of the problem of analyzing indirect discourse in general—which includes ascriptions of intensional attitudes and claims about what other people have said—and the challenges that quotational approaches to that kind of discourse have to face (the quotational approach takes the believer/speaker to be related to a *sentence* rather than to its meaning or the proposition expressed by it).[51] In fact, our truth-conditions for ascriptions of fine-grained information-possession bear some resemblance to candidate truth-conditions for

---

[50] There is a bulky literature on what *makes it the case* that vehicles have the content they have, and in virtue of what facts are them representational—see Cummins (1989), Ramsey (2001), Shea (2018) for some examples. But we won't try to tackle that issue here. We are taking it for granted that representation *does* occur, without trying to determine *why* and *how* more exactly it occurs. More important to our theoretical goals, there must be a systematic way of assigning coarse-grained information to our vehicles, in such a way as to respect entailment relations among them. And a valuation function that maps the vehicles that constitute fine-grained information into coarse-grained information must do justice to the set-theoretic connections (subset-relations) between the relevant bits of coarse-grained information, too. For example, the coarse-grained information carried by a vehicle that depicts an object *a* as being *F* and *G* (for example, an object as being round and red) must be a *subset* of the coarse-grained information carried by a simpler vehicle that depicts *a* as being *F* while being silent regarding whether *a* is *G*.

[51] See for example the discussion in Quine (1960, Chapter VI), Davidson (1968) and Lepore and Loewer (1989).

ascriptions of attitudes that refer to *both*, linguistic expressions and their semantic values.[52]

§4.2 Where the distinction between fine-grained and coarse-grained information is further explicated (continued)

One idea is that Euclid was thinking with an internal *translation* of our English sentence into his own language, and that it is some such translation that constituted his fine-grained information. In order to make space for other kinds of representational vehicles than natural language sentences (say, images, maps, diagrams), however, we would have to either broaden the notion of translation to include non-linguistic representations among its *relata*, or rather put things in terms of a more general similarity relation among representational vehicles, where the relation of *being a translation of* is but an instance of that similarity relation.[53]

We can ground the relevant similarity relation on a number of different properties of/relations between representational vehicles, such as how they are combined together to represent the world (think, for example, of an isomorphism between the grammatical form of sentences and the way icons are spatially arranged in an iconic representation), the inferential roles that they play within the inquirer's cognitive system, and their semantic properties and relations (for example, sameness of intension). Still other properties and relations might be relevant for that purpose.

We adopt that second option, thus putting aside the problem of determining exactly when a representational vehicle constitutes a translation of the other. The idea is that, where '*p*' is taken from an ascription of fine-grained information such as 'she possesses the fine-grained information that *p*', there will be some similarity class *f*('*p*') of representational vehicles that are similar to '*p*' under the relevant dimensions.

Importantly, all the members of such a similarity class must carry the same coarse-grained information. That is, where both vehicles *x* and *y* belong to *f*('*p*'), it follows that $v(x) = v(y)$. The function *v* is again a valuation function, which maps representational vehicles into sets of possible worlds. $v(x)$ is the set of possible worlds where vehicle *x* holds true or is accurate. The domain of our valuation function now

---

[52] See for example the proposal put forward by Richard (1990). Note, however, that coarse-grained contents in Richard's framework are the so-called 'Russellian propositions', as opposed to our sets of possible worlds.

[53] A fair amount of research has already been done concerning the translation of certain kinds of representations into/from other kinds of representations in cognitive science and artificial intelligence. See for example Larkin and Simon (1987) for the 'informational equivalence' among diagrams and sentences, Coppin et. al. (2016) for the translation of iconic representations into text, and Saha, et. al. (2021) for the idea of translating images into maps. On the idea that maps are vehicles of thought, or that they have the properties that a vehicle needs to have in order to play that role, see Camp (2007).

includes not only sentences of the language we are *using* to ascribe information possession (here '*p*'), but also the representational vehicles that are similar to those sentences under the relevant dimensions.

The vehicle '*p*' itself must belong to *f*('*p*'), of course, for identity is but a special case of the relevant similarity relation (every vehicle is similar to itself). It follows from this that any vehicle that belongs to *f*('*p*') carries the same coarse-grained information as the one that is carried by '*p*'. In other words, $v(x) = v(\text{'}p\text{'})$, for any $x \in f(\text{'}p\text{'})$.

We can now attempt to state the truth-conditions for ascriptions of fine-grained information possession. We say that 'the inquirer has the fine-grained information that *p*' is true at *w* if and only if there is a representational vehicle *x* that belongs to *f*('*p*') such that, at *w*, *x* is part of the inquirer's representation of how *w* is like. Having the fine-grained information that *p* is a matter of one's having incorporated some representational vehicle similar to '*p*' (similar under the relevant dimensions) into one's internal model of reality.

We won't try to analyse that further, or break the sentence '*x* is part of the inquirer's representation of how the world is like' down to more basic conditions. Minimally, however, that *x* is part of the inquirer's representation of reality at *w* means that (a) the inquirer thinks of reality as being such that *p through* that vehicle *x* (or *x* is a means by which the inquirer thinks of reality as being such that *p*), and (b) $i_w \subseteq v(x)$, where i_w is again the inquirer's body of coarse-grained information at *w*.[54]

Explicating what it is for a representational vehicle to be part of one's representation of reality further would require taking a number of different considerations into account. Notice, for example, that when the inquirer deploys a certain vehicle *x* to represent the world as being a certain way *in her imagination* (say, *x* is a percept-like representation), that vehicle doesn't thereby count as part of the inquirer's representation of reality. But such a vehicle *could* be part of her representation of reality, for example, if it were constitutive of the inquirer's *perception* of her surrounding environment (similarly, contrast *supposition* to *belief*). There is a difference between perceptually representing *w* as being such that *p* and imagining *w* as being such that *p*—and that difference is but one of the data points to be taken into account by any attempt to formulate criteria for when a vehicle belongs to the inquirer's representation of how the world is like.

Be that as it may, we have now an entailment relation between the truth-conditions for ascriptions of fine-grained information to ascriptions of coarse-grained information, but not the other way around, as desired.

---

[54] We can also relativize such bodies of information to different fragments in the inquirer's mind, so that i_w might differ relative to two different fragments.

Suppose that 'the inquirer has the fine-grained information that *p*' is true at *w*. Then there is a vehicle *x* such that *x* is part of the inquirer's representation of reality at *w* and $x \in v(\text{'}p\text{'})$. It then follows from this that *x* and '*p*' carry the same coarse-grained information, that is, $v(x) = v(\text{'}p\text{'})$. By condition (b) from above, it follows furthermore that the inquirer's coarse-grained information i_w is such that $i_w \subseteq v(x)$. Finally, it follows from the two previous consequences that $i_w \subseteq v(\text{'}p\text{'})$. But that is just the truth-condition of 'the inquirer has the coarse-grained information that *p*' at *w*. So the fact that the inquirer has the fine-grained information that *p* entails that the inquirer has the coarse-grained information that *p*.

The entailment in the other direction doesn't hold, however, and that is because the fact that $i_w \subseteq v(\text{'}p\text{'})$ does not entail that there is a representational vehicle *x* that belongs to $f(\text{'}p\text{'})$ such that *x* is part of the inquirer's representation of how *w* is like. The fact that the inquirer's coarse-grained information is a subset of the set of worlds where '*p*' is true doesn't guarantee that a representational vehicle that is similar to '*p*' (similar under the relevant dimensions) is part of the inquirer's internal model of reality. Her representation of reality might be constituted by *other* vehicles, which together entail members of $f(\text{'}p\text{'})$, without being constituted by any member of $f(\text{'}p\text{'})$ in particular.

For example, suppose that it is part of your fine-grained information (say, because you came to know these things) that *Cat Power is a singer* and that *Cat Power is Chan Marshall*. So your coarse-grained information is constituted by possible worlds where both '*Cat Power is a singer*' and '*Cat Power is Chan Marshall*' are true, and both a vehicle that is similar to the former sentence and one that is similar to the latter one are constitutive of your representation of how the world is like.

Since the sentence '*Chan Marshall is a singer*' is also true in all of those worlds, you count as possessing the coarse-grained information that *Chan Marshall is a singer*. But that is not part of your fine-grained information yet, we may assume, because you haven't yet 'put two-and-two together' in the following sense: you have the two bits of fine-grained information that *Cat Power is a singer* and that *Cat Power is Chan Marshall*, respectively, but you haven't yet inferred from there that *Chan Marshall is a singer*. So there isn't yet a vehicle *x* that belongs to the class $f(\text{'}$*Chan Marshall is a singer*$\text{'})$ such that *x* is integrated into your internal model of reality. And adding such a vehicle *x* to it is just what you do when you finally infer that conclusion from those premises. Before the inference, your representation of reality was constituted by vehicles that together entail *x* (they are true or accurate only if *x* is true or accurate), without itself being constituted by *x*—and similarly with respect to other members of $f(\text{'}$*Chan Marshall is a singer*$\text{'})$. We are assuming for the sake of illustration here that '*Cat Power is a singer*' is *not* one of the members of $f(\text{'}$*Chan Marshall is a singer*$\text{'})$.

Similar points also serve to show that ascriptions of fine-grained information do not have the 'closure' property mentioned above. The fact that the inquirer possesses the fine-grained information that *p*, and that *p* entails that *q*, does not guarantee that the inquirer possesses the fine-grained information that *q* (though it guarantees that the inquirer has the coarse-grained information that *q*).

§4.3 Where a new model of inquiry is fleshed out

That is, then, our way of addressing the puzzle of armchair inquiry presented in §3.3.

Suppose an inquirer settles her question *Q* by validly deducing a maximal answer to it from the information possessed by her—the answer that *p*. In a sense, the inquirer already had a complete answer to *Q* even before coming to settle *Q* in this way (answer as coarse-grained information). For, by assumption, *v*('*p*') ∈ *Q*, or the set of possible worlds where '*p*' is true is a complete answer to *Q*. Since the inquirer's deduction was valid, then, her coarse-grained information was a subset of *v*('*p*') already before she performed that inference. That is, the inquirer's coarse-grained information was already a member of *Q* (for *Q* is again the set of its complete answers/the downward closed set of its maximal answers).

According to the truth-conditions from the previous section, we now say that our inquirer has the coarse-grained information that *p* both *before* and *after* making the deductive inference. That is the sense in which the deduced answer is *not new* information with respect to that inquirer. She comes to acquire the fine-grained information that *p* after performing the deduction, however. And that is the sense in which the deduced answer is *new* information to her. The process of deduction allows our inquirer to incorporate a new vehicle *x* that belongs to the similarity class *f*('*p*') into her representation of reality.

In §3.3, we said we would explore the limitations of the possible-worlds model of inquiry not only from the angle of the problem of the informativity of deduction, but also from the angle of our taxonomy of questions.

That taxonomy, remember, featured three types of questions, relative to any inquirer at any time: the questions that the inquirer is already in a position to settle (I^+S^+), those that the inquirer has enough information to settle but is not yet in a position to settle (I^+S^-), and those that the inquirer doesn't yet have enough information to settle (I^-). The problem with the possible-worlds model of inquiry was that it lumped together questions of type (I^+S^+) and questions of type (I^+S^-), despite the fact that the inquirer is in a different epistemic standing with respect to those two types of question. Questions of type (I^+S^+) are questions that the inquirer can settle through armchair inquiry. Questions of type (I^+S^-) are not yet questions that the inquirer can settle

through armchair inquiry: even though she has enough information to settle them, she still lacks the cognitive skills to do so.

Now our proposal about the informational yield of deduction from above can be used to fetch a new model of inquiry that is also in a broad sense a possible-worlds model. In addition to that, however, it also represents the bits of *fine-grained* information that are possessed by the inquirer at each possible world. Questions of type (I^+S^+) and questions of type (I^+S^-) can be distinguished within this model—an improvement over the old model.

The new model continues to represent questions as we did before: they are the sets of all their complete answers, construed as sets of possible worlds. Besides telling us which possible worlds are part of the inquirer's body of coarse-grained information, however, its frames also tell us which vehicles are part of her representation of reality.

More formally, the new frames deliver (for every possible world w) not only i_w, or the set of worlds that are compatible with the inquirer's fine-grained information at w, but also a set of representational vehicles r_w, which stand for the ways in which the inquirer represents w as being when she is in w. The members of such a set r_w are sentences in a language of our choice—its members stand for (as opposed being identical to) the vehicles that can be individuated from the inquirer's representation of reality. So i_w captures *what* the world w must be like from the inquirer's own perspective, and r_w captures *how* the inquirer represents w as being that way.

As we saw before (§3.1), the old possible-worlds model is able to track progress in information-gathering inquiry via updates on the set of possible worlds i_w that constitutes the inquirer's coarse-grained information at w. As the inquirer learns new facts, that set is chopped down into ever smaller sets of possibilities (in formal terms: the accessibility relation of the updated frame is less inclusive than the one from the original frame).

But now we can also think of deductive updates that keep the inquirer's coarse-grained information i_w *the same*, and the only thing that changes after the update are the members of r_w, that is, the vehicles that constitute the inquirer's representation of how w is like. Instead of outputting a new set of possible worlds, these update operations output new representational vehicles, corresponding to increments in the inquirer's internal model of reality—novel depictions of the same facts.

We can put it roughly as follows. There is the set of possibilities that are open to the inquirer, and then there is the set of vehicles that the inquirer has *read off* of that set of possibilities. In the new model of inquiry, updates might modify the latter without modifying the former. Progress in information-gathering inquiry is tracked by the elimination of possible worlds, progress in armchair inquiry is tracked by the addition of representational vehicles.

The update that tracks progress in information-gathering inquiry is an update on a *new bit of information*, and the update that tracks progress of armchair inquiry is an update on a *new inference* performed by the inquirer. What is new after the latter kind of update are only new members that are added to the inquirer's set of vehicles r_w.[55]

So the kinds of inferences that the inquirer is able to perform at different stages of inquiry must also be represented within this model. We add to its frames, then, a set of *inference rules* that the inquirer can apply to the vehicles that are already part of her representation of reality, so as to enrich that representation by adding new vehicles to it. These rules take a set of representational vehicles as input, and they output a new set of representational vehicles.[56] When the inquirer applies an inference rule to r_w, then, the rule outputs a new set r^*_w of vehicles, standing for the inquirer's enriched representation of reality after she has deduced an answer to a question she was inquiring into before. (The inquirer can also be represented as having different inferential rules available to her at different worlds, so as to make room for update operations that track the inquirer's learning of new inferential skills).

Questions of type (I^+S^-) and questions of type (I^+S^+) *can* be distinguished within this model. Where the inquirer has information that is enough to settle a question Q, but lacks the cognitive skills that are needed to settle it on the basis of information possessed by her at world w, her situation is represented by a model with the following features: (a) the inquirer's coarse-grained information i_w is a member of Q, but (b) there is no inferential rule available to the inquirer at w such that, when applied to the members of r_w, it outputs a new set of vehicles r^*_w one of whose members belongs to the class $f(\text{`}p\text{'})$. That represents the fact that the inquirer cannot get a particular kind of fine-grained information out the information she already has. In contrast, when it comes to a question of type (I^+S^+), the model that represents the inquirer's situation satisfies (a) while also satisfying the *denial* of (b): there *is* an inferential rule…

Call that the *new model of inquiry*. It tries to capture both, the inquirer's coarse-grained and her fine-grained information at each possible world. And it allows us to represent at least some of the cognitive skills that the inquirer can deploy in her attempt to settle a question from the armchair, namely, her inferential skills.

No doubt this whole framework will be much more complex than the one that deals only with coarse-grained information, leaving representational vehicles and

---

[55] There are already systems of dynamic epistemic logic that implement just this idea—see again van Benthem and Velázquez-Quesada (2010).

[56] These rules will have to abide by a number of constraints. In particular, where the inference rules are rules whereby reasoners *increase* their knowledge via deduction, the set of vehicles that constitutes the input to the rule must be a proper subset of the set of vehicles that constitutes its output, so that no representational vehicles is lost after the inference. We won't complicate things even more here by making room for the representation of information loss.

cognitive skills aside. But here we have at least the barebones of a formal implementation of our solution to the puzzle of armchair inquiry from above, in such a way as to also distinguish questions of type (I^+S^+) from questions of type (I^+S^-).

§4.4 That we can now build a system of instrumental norms of inquiry on the basis of our theoretical constructs

Our new model of inquiry, as well as the idea on which it is based, offers us a picture of what armchair and information-gathering inquiry in general are about. The picture can be summarized as follows. Armchair inquiry trades on *vehicle* manipulation without necessarily altering *content*. In contrast, information-gathering inquiry trades on directly altering content. ('Content' here stands for coarse-grained information.)

As we saw, there is a sense in which armchair inquiry allows the inquirer to acquire *new* information, even in the cases of deduction discussed above—but it is *fine-grained*, not coarse-grained information. Fine-grained information is representational vehicle *with* content, not the content itself. So the two kinds of inquiry are still importantly different.

This is not again to deny that the two types of inquiry often occur as parts of a single, larger project of inquiry. When doing empirical research, we collect coarse-grained information, which allows us to eliminate previously uneliminated possibilities. But then we still have to reason from the newly collected information (together with the information we already had before) to settle some of the questions we were inquiring into. Even empirical research involves armchair inquiry. Heavy use of it is made not only in the areas of inquiry that are typically labelled as 'armchair' areas of inquiry (philosophy, mathematics), but also in the natural sciences in general.

In §1.3, we drew a contrast between questions that call for information-gathering and questions that can be settled through armchair inquiry, as follows: given the information that is already available to an inquirer at any given point in time, there are questions such that her being able to settle them at that time strictly depends on her gathering *more information* from that time onwards, on the one hand, and questions that she is able to settle at that time *regardless* of whether she gathers more information from that time onwards, on the other.

And now we make that more precise by pointing out that what is meant by 'more information' there is more *coarse-grained* information. Some questions are such that the inquirer needs to gather new coarse-grained information first in order to come to be in a position to settle them (only *then* will she be in that position). And there are questions such that the acquisition of new coarse-grained information is not strictly needed in order for the inquirer to come to be in a position to settle them (she is *already* in a position to settle them), even if she still needs to make stepwise, reasoned

out additions of representations to her internal model of reality, until she gets to a point where she settles her question. That she needed new fine-grained information, or new representations of the same facts, doesn't mean that she needed new coarse-grained information in order to settle her question. The *new* information that she needed could be 'gathered' by her from the comfort of her armchair, with her eyes shut (no need for smelling, tasting, hearing or touching things either).

In the new model of inquiry, the acquisition of such new information is tracked, again, by updates on the set of vehicles that constitute the inquirer's representation of reality, which allow her coarse-grained information to remain intact before and after the update (no new bit of coarse-grained information is needed for such an update).

There should be norms that tell the inquirer to engage in information-gathering inquiry, as well as norms that tell him to engage in armchair inquiry. Compliance to such norms is supposed to further the goals of inquiry—crucially, the goal of settling a question.

The norms that call for information-gathering inquiry will tell the inquirer to look for new coarse-grained information, so that the space of possibilities that are open to him gets chopped down into a smaller set, in such a way that he comes closer to a situation where he is in a position to settle the question he is inquiring into. In contrast, the norms that call for armchair inquiry will tell the inquirer to expand his repertoire of fine-grained information, again in such a way as to at least come closer to a situation where he is in a position to settle the question he is inquiring into.

More specific norms that call for information-gathering inquiry (norms that will tell the inquirer more precisely what to do) can be seen as all stemming from or as being grounded on the following norm of inquiry, where 'GG' stands for 'Go Gather':

(GG) If one wants to settle *Q*, but one's coarse-grained information is not yet enough to settle *Q*, then one should see to it that one gathers new information that bears on *Q*.

Information that bears on *Q* is information that constitutes at least a partial answer to *Q*.

Suppose, for example, that you want to settle the question of *who called you* (on your cell phone). The phone number of the person who called you is unknown to you, and it is not the phone number of any of the contacts you have saved on my phone. Your coarse-grained information, say, as determined by what you know or by what your evidence is, is not yet enough to settle that question—it is not any of the complete answers to that question (that would be some set of possible worlds such that *c called you* is true at all the members of that set, for some person *c*).

Given that much, (GG) tells you to gather more information, as long as you still want to settle the question of *who called you*, until you end up with enough of it to settle that question (or rather you cease to want to know what its answer is). And that, we might suppose, is just what you do: you call the number back to find out who it was, or you send them a message, etc. If none of these actions provide you with the information you need, however, you leave the issue aside and stop wanting to settle *who called you*. If we read the instrumental 'should' of (GG) as having wide-scope—which is the reading we adopt here—then these are both ways in which you abide by (GG). Since you are acting rationally when you act in either of these ways, it seems that our norm gets things just right.

To be more precise, (GG) has the logical form $S(\varphi \wedge \psi \supset \sigma)$, where '$S$' is the *should* operator, '$\supset$' is the material conditional, '$\wedge$' is the conjunction sign, 'φ' is a placeholder for 'one wants to settle Q', 'ψ' is a placeholder for 'one's coarse-grained information is not yet enough to settle Q', and 'σ' is a placeholder for 'one sees to it that one gathers new information that bears on Q'.

(GG) tells inquirers to go gather more information when certain conditions are satisfied. So it is a norm that directs inquirers to engage in information-gathering inquiry. Other norms will direct inquirers to engage in armchair inquiry. Consider for example the slightly more complicated norm, where 'GF' stands for 'Go Figure':

(GF) If one wants to settle Q, and one's coarse-grained information i is enough to settle Q, then one should see to it that one competently deduces some fine-grained information x such that $v(x) = a$, where a is the maximal complete answer to Q such that $i \subseteq a$.

Whereas the Go Gather Norm (GG) tells inquirers to collect more coarse-grained information when certain conditions are satisfied, the Go Figure Norm (GF) tells them to process the fine-grained information they already have when some alternative conditions are satisfied.

(GF) directs inquirers to figure things out by exploiting the information they already have at their disposal through deductive reasoning. For, if an inquirer satisfies the conditions described in the antecedent of (GF), she can in principle get what she wants by means of sheer reasoning. In a sense, again, she already has the answer she is looking for, implicit in her total body of information as it may be.

The formulation of (GF) is a bit cumbersome. The function v, remember, takes a representational vehicle as input, and it outputs the set of worlds where that vehicle is true. So $v(x)$ is the set of worlds where x is true—it is the body of coarse-grained

information carried by vehicle *x*. The phrase 'some fine-grained information *x* such that $v(x) = a$, where *a* is the one maximal complete answer to *Q* such that $i \subseteq a$' refers to *what kind of vehicle* the inquirer is told to deduce by (GF).

To break that down and make it clearer, notice first that *if* the inquirer's coarse-grained information *i* is enough to settle *Q* (as stated in the antecedent of the conditional embedded in (GF)), then *there is* a maximal complete answer *a* to *Q* such that *i* is a subset of that maximal complete answer, or $i \subseteq a$. That maximal complete answer *a* is a bit of coarse-grained information that is carried by representational vehicles of a certain sort—which include natural language sentences through which speakers respond to *Q* out loud or in paper. It is some such representational vehicle that (GF) tells inquirers to add, via deduction, to their body of fine-grained information or their picture of reality.

In the new model of inquiry, improvements in inquiry brought about by compliance to (GG) are tracked by updates in function *i*, which tells us which possibilities are compatible with the inquirer's information at different possible worlds. In contrast, improvements in inquiry brought about by compliance to (GF) are tracked by updates in function *r*, which tells us which vehicles are part of the inquirer's representation of reality at which possible worlds.

The Go Gather (GG) and the Go Figure (GF) norms can be seen as building blocks of a comprehensive system of instrumental norms of inquiry.[57] We can also see them as more specific instances of a general 'Zetetic Instrumental Principle' (ZIP), named this way after Friedman (2020, p. 503), according to which one should take the means to figure out the answers to the questions one wants to know the answer to (though notice that Friedman uses 'ought', and that she gives it a narrow-scope reading —in contrast to our use of 'should' with a wide-scope reading). Given that an inquirer wants to know *Q*, for some question *Q*, *at least one* and *at most one* of those norms will 'fire', or it will apply to that inquirer's case, in the following sense: the inquirer needs to eventually *do* something in order not to violate the corresponding norm.

Suppose that the inquirer wants to settle *Q*. Either his coarse-grained information *i* is enough to settle *Q*, or it isn't. If the former, then the Go Figure Norm calls for action on the part of that inquirer. If the latter, then the Go Gather Norm does. There is time to do information-gathering inquiry and time to do armchair inquiry—but one thing at a time. There is time to collect data and time to reason from that data. One can only draw inferences from the data after one has collected it. As long as the inquirer wants to settle the question, no matter how his total coarse-grained information

---

[57] For an overview of the norms of inquiry that have been mostly discussed in the recent literature, see Haziza (2023).

stands with respect to it, one or the other of these norms will tell us what that inquirer should do at that time.

We continue our investigation by submitting these norms to critical scrutiny. They also need to be refined, in that some of their important parameters (if they are to be true) are hidden from our formulations above. In the next chapter, we occupy ourselves with the Go Gather Norm. Chapter 6 turns to the Go Figure Norm, clarifies the nature of ‘should’ in both of these norms, and fleshes out further norms of inquiry, which are prohibitive norms.

Chapter 5

§5.1 That there are reasons to reject the Go Gather Norm

The Go Gather norm (GG) tells the inquirer to gather new information that bears on the question she wants to settle, in case that inquirer doesn't yet have a complete answer to it (in the sense of coarse-grained information). As long as the inquirer wants to settle her question, (GG) keeps telling her to collect data, until she ends up with information that is enough to settle her question.

That sounds plausible, again. If you want to settle the question of *what the boiling point of gold is*, for example, and the information you have isn't yet enough to settle that question, then it looks like you should read chemistry books, or maybe try to measure the temperature at which gold starts melting yourself (if you're in a situation where you cannot rely on any testimony concerning such measurements), etc. And that is exactly the kind of thing that (GG) says you should do. That looks indeed like a straightforward case of means-ends rationality: given the goal that you have (you want to settle *Q*), and given the kind of situation you find yourself in (not having enough information to settle *Q* yet), you should take the means to your goal in that situation (gather more information).

In interpreting the notion of *information that is enough to settle a question* in the way we did in §2.3, however, one might well dispute the truth of (GG). Remember how we have explicated that notion. The inquirer's body of coarse-grained information is some set of possible worlds. The answers to a question are also sets of possible worlds. The question itself is the downward closed set of its maximal answers. And we have assumed that the inquirer's information is enough to settle a question *Q* just in case her body of coarse-grained information is a *member* of *Q*. We can also put it as follows: the inquirer's information is enough to settle *Q* just in case the information possessed by her *entails* one of the maximal answers to *Q*.

And it might be objected that some bodies of information are enough to settle a given question *Q* even though they do not so much as entail any maximal answer to *Q* —maybe they only make one of those maximal answers highly probable. Otherwise, (GG) becomes *too strong* a norm. For we don't always need to see to it that our bodies of information *entail* the maximal answers to the questions we want to settle.

The objector might be thinking here, for example, of the use of abductive reasoning to settle questions. Suppose the inquirer wants to settle *Q,* but her coarse-grained information isn't yet enough to settle *Q* in our sense. Instead of collecting more

information, she rather comes up with the best explanation she can muster of certain bits of her total body of information. And the *explanans* she comes up with is one of *Q*'s maximal answers. It seems that she thereby settles her question, and she does so from the armchair, in that she comes to form a belief in one of the question's maximal answers by reasoning from the information she already has. Indeed, assuming that the relevant kind of reasoning is reliable, then what we have said so far doesn't *preclude* that she has settled *Q* in the sense of §2.1, where *settling a question* was taken to be some form of epistemic success, at a minimum reliably formed belief in one of the question's answers.

But if the notion of *information that is enough to settle a question* is interpreted in the way we did, then (GG) says that our inquirer should have gathered more information, given our assumptions. Since that is not what she did, however, she has violated (GG) under that interpretation. Instead of gathering more information until she had enough of it to *entail* one of the maximal answers to *Q*, she has inferred one of the question's answers by means of a reliable but fallible method of reasoning. But it seems that she hasn't thereby violated any norm of inquiry. So (GG) must be false under our favored interpretation of the notion of *information that is enough to settle a question*.

Consider the objection also from the angle our inquirer-relative typology of questions. One of our classes there was the class of questions of type (I^-) for an inquirer at a particular time, that is, any question *Q* such that that inquirer doesn't have enough information to settle *Q* at that time yet. According to the framework we have adopted, that is just a question *Q* such that the inquirer's coarse-grained information doesn't constitute an answer to or is not yet a member of *Q* (no maximal complete answer to *Q* is entailed by the inquirer's information). Now perhaps the class of questions of type (I^-) is just *too big* for any typical inquirer. The contention is that *only rarely* do we have information that is enough to settle a question in the present sense.

For the sake of illustration, take run-of-the-mill cases of inquiry. Farmer Joe is worried about his crop. As he wakes up one morning, he asks himself: *Did it rain last night*? He gets up and goes to the front door. He opens it up and takes a good look at the lawn in front of his house. He immediately realizes that the grass is wet. He infers on that basis that *it rained last night*—and that is the happy end of Joe's inquiry. For, at that point, Joe is not inquiring into *whether it rained last night* anymore, since he has already settled on an answer to his question (or so it seems). But does the information possessed by Joe—say, as determined by what he knows or what his evidence is—*entail* that *it rained last night*?

If not, then Joe has closed his inquiry without having information that was enough to settle his question under our favored reading, and he has failed to abide by

(GG) under that same reading. For he did not collect more information until he had enough of it to entail that *it rained last night*. Instead of doing that, Joe went ahead and inferred one of the question's answers on the basis of less than conclusive information. And now it is as if his coarse-grained information were to get chopped down into a smaller set (no more possible worlds where it *didn't* rain last night) from the comfort of his armchair, namely, by means of that risky inference he made. By sheer armchair thinking, he is ruling out possibilities that *could* be the world he is in, as far as the information possessed by him goes, which includes the information that *his grass is wet* and, let us suppose, the information that *rain causes the grass to be wet*. Joe is aware that other things could have made the grass wet, of course, but the hypothesis of rain overnight is the best explanation he can find for it.

Cases such as this one multiply easily. It seems that scientists very often draw such risky inferences, for example. It seems, that is, that they often deliver answers to questions (such as the question of *whether humans and dinosaurs coexist*) even though the information possessed by them is not enough to settle those questions in our sense.

§5.2 Where we explore ways to salvage our framework

We can offer a number of different responses to that worry against (GG) under our favorite interpretation of 'information that is enough to settle a question'. Such responses are attempts to protect the integrity of our whole theoretical framework. We start developing them in the present section and the next (§5.3), and we bring them to bear on (GG) in §5.4.

The first relevant consideration here is the following. In many cases where inquiry is resumed through abductive reasoning and the like, the inquirer *does* settle some question, and one that she did have enough information to settle in our sense—though it is not exactly the question that we originally described her as inquiring into. It is rather some other question, and one that it was *easier* for her to settle.

To follow up on a strategy that has already been hinted at in §2.2, we can think of a probabilistic rendering $p(Q)$ of a question Q, where $p(Q)$ is something like: *Which of the maximal complete answers to Q is more probable than the others*? Where Q is the polar question *Did the universe start with a Big Bang?*, for example, $p(Q)$ would be the question *Which of these is more probable: that the universe started with a Big Bang or that it didn't?*

Let us continue with that example, and interpret ‘probable’ as expressing evidential probability, or probability relative to a body of evidence.[58] Scientists who are in possession of evidence that makes it highly probable that *the universe started with a Big Bang* may be in possession of coarse-grained information that constitutes an answer to *p*(*Q*), such as the information that *it is more probable that the universe started with a Big Bang (more probable on the relevant evidence),* even though they are not in possession of coarse-grained information that constitutes an answer to *Q*, such as the information that *the universe started with a Big Bang*. They are sufficiently informed about *what their evidence is* and *how it bears on the issue*, without being sufficiently informed about *how the world is*.

Inquirers often use non-deductive reasoning to defeasibly infer an answer to a question *Q* on the basis of their evidence and background assumptions. If they do that properly, they show sensitivity to how their evidence bears on the answers to *Q.* In particular, they thereby show sensitivity to how probable some of those answers are conditional on their evidence, in comparison to the other answers. In many cases of this sort, their coarse-grained information is a member of *p*(*Q*) even though it is not a member of *Q*: they have the coarse-grained information that one of the answers to *Q* is more probable than the others, even though they do not have the coarse-grained information that a_1, neither do they have the coarse-grained information that a_2, …, neither do they have the coarse-grained information that a_n, where a_1, …, a_n are all the possible maximal answers to *Q*.

(Keeping these points in mind, we might even interpret an inquirer’s reply to the utterance of an interrogative whose semantic value is *Q*—say, in public communication —as conveying the information that one of the answers to *Q* is much more probable than the others, despite the fact that her reply consists of the utterance of a declarative sentence whose semantic value is one of the maximal complete answers to *Q*.)

A second consideration that bears on the present concern is the following. Felicitous use of phrases such as ‘settle a question’ and ‘information that is enough to settle a question’ can be made in *loose speech*, even in situations where the sentences in which those phrases occur are literally speaking false.[59]

Felicitous uses of ‘the inquirer has settled the question’ can be made, say, to describe situations where the inquirer has rationally adopted some attitude of *acceptance* towards one of the maximal answers to a question she has inquired into—an attitude the holding of which involves her acting as if that answer is true, at least for

---

[58] See Williamson (2000, Ch. 10) and Eder (2019) for two alternative ways of understanding the notion of evidential probability.

[59] See Bach (2001), Carter (2021) for some important properties of loose talk. See also Moss (2019) for how it might bear on the semantics of ascriptions of belief.

some practical purposes.[60] And there is a corresponding use of 'information that is enough to settle a question', namely, one that applies to situations where the inquirer possesses information that is enough to make it rational for that inquirer to accept one of the maximal answers to the question, or to act as if that answer is true, at least for certain practical purposes.

Concerning the case of Farmer Joe from the previous section, for example, we might felicitously utter 'Farmer Joe's information is enough to settle the question of *whether it rained last night*' to convey the information that it is rational for him to accept that *it rained last night* and carry on with his affairs under the assumption that that is the case. Such uses are felicitous, in the sense that speakers thereby get to share the information that they wanted to share with their hearers in an effective manner (without cluttering their minds with more complicated constructions). But the target utterances are still *literally speaking* false.

We do not need to go deeper into the details of such loose uses of the target expressions. We do not need, for example, to try and determine whether the notion of rationality that is presupposed there is 'non-epistemic', or whether it is 'epistemic' though it is 'pragmatically encroached', etc.[61] The point is that the acknowledgment of felicitous uses of 'settle a question' and 'information that is enough to settle a question' is compatible with the claim that those expressions have stricter standards of application, such as those imposed by the framework we have developed here.

In loose speech, we can felicitously utter tokens of the type 'The inquirer has information that is enough to settle *Q*' even though the inquirer's information is *not* one of the members of *Q*, or it does not entail any of *Q*'s maximal answers, making one of them highly likely to be true as it may. Taken literally, however, 'information that is enough to settle *Q*' refers to a body of information relative to which *Q* is *closed*: no more than one of *Q*'s maximal answers can be true given that body of information. The question must, after all, be *settled* relative to the relevant body of information, and to be settled is to be *completely* settled.

So those utterances in cases such as Farmer Joe's are strictly speaking false. But they may still be felicitous, in that we thereby get to share the information that we wanted to share with our hearers. (Compare: in uttering the literally false sentence 'I

---

[60] Thanks to Andy Müller and Arianna Falbo for discussion on this. Notice that there at least two available readings of 'settle' when it is used as a transitive verb in this manuscript, though it takes the same object in both cases, namely, a term that refers to a question: (a) with a term that refers to an inquirer as its subject, (b) with a term that refers to a body of information as its subject.

[61] For some discussion in the literature that bears on such issues, see for example Fantl and McGrath (2002), Schroeder (2012)

arrived at 5:00 p.m.', say, in a situation where you arrived at 5:02 p.m., you get to convey the information that *you arrived at a time close to 5:00 p.m.* to your hearers).[62]

Let *Q* be a polar question with the two maximal answers *yes* = {w_1, w_2, w_3, w_4} and *no* = {w_5, w_6, w_7, w_8}. Suppose Ana's body of coarse-grained information is i_A = {w_1, w_2, w_3}, and Bob's is i_B = {w_1, w_2, w_3, w_5}. Ana's information is enough to settle *Q*, but Bob's isn't. For, whereas Ana's body of information is constituted by *yes*-worlds only, Bob's is constituted by a mix of *yes*- and *no*-worlds. The *no*-possibility w_5 is still there, bothering Bob. Since he didn't learn things to rule w_5 out, as it happened to Ana, he doesn't yet have information that is enough to settle *Q*, like Ana does. Relative to Ana's body of information, the question is *closed*. Relative to Bob's body of information, however, it is still an *open* question.

But we can easily conceive of further details that would make it felicitous for us to say that Bob's information, and not only Ana's information, is enough to settle *Q*: it is *very improbable* that w_5 is the world that Bob is in, relative to the evidence possessed by him.[63] And so it is rational for Bob to accept the *yes* answer to question *Q*, to act as if it is true for certain practical purposes, to ignore the possibility that the *no* answer is true where he has to make decisions that hinge on what the answer to *Q* is (the world around him makes pressure on him to make up his mind and act fast), etc.

That we can successfully impart all this information by simply uttering 'Bob's information is enough to settle *Q*' does not mean that Bob's information *is* enough to settle *Q* (even Bob himself might acknowledge that much).

§5.3 Where we explore ways to salvage our framework (continued)

The two considerations from the previous section are meant to appease the worry that the notion of *information that is enough to settle a question* from our theoretical

---

[62] Some of the good reasons for thinking that such a sentence is false in the envisioned kind of situation are presented and discussed by Lasersohn (1999).

[63] Suppose that *pr* is a probability function that measures evidential support, and let its distribution over our total set of possibilities be partly characterized by $pr(\{w_1\}) = pr(\{w_2\}) = pr(\{w_3\}) = 0.2$ and $pr(\{w_5\}) = 0.05$ (the values it assigns to the other possible worlds need not occupy us here). It follows that the *yes*-answer is more likely, in fact *much* more likely, than the *no* answer conditional on Bob's total body of information i_B. This is computed in the usual way as follows:

$pr(yes \mid i_B) = pr(yes \cap i_B) / pr(i_B)$,
$pr(yes \mid i_B) = 0.6/0.65$,
$pr(yes \mid i_B) \approx 0.92$.

In contrast, $pr(yes \mid i_A) = 1$, where i_A is Ana's body of information. Probabilistically speaking, it cannot get better than that for Ana (not by learning *more* things without loosing any previously gathered information). But things can definitely get better for Bob: he may acquire new information that will bump the probability of *yes* up to 0.95, 0.98, or 1.

framework is too demanding. They tell us not to worry so much about it being too rare or too difficult for an inquirer to possess information that is enough to settle a question —for there is also *this other fact* (a different fact is adduced to by each consideration), and *this other fact* should put us at ease and allow us to come to terms with the assumptions of our framework.

In the case of the first consideration, *this other fact* is the fact that there is often some other question (other than *Q*), such as the question of *which of the answers to Q is more probable than the others*, such that the inquirer's information is enough to settle it. In the case of the second consideration, *this other fact* is the fact that our utterances of expressions such as 'the inquirer has information that is enough to settle the question' are often *felicitous*, though literally speaking false—where felicity is a function of how well the speaker can convey the information he intends to convey to his audience by uttering those expressions.

The additional facts that those two considerations allude to are supposed to prevent the balance from tipping to 'no' when we weigh the pros against the cons of our framework and try to decide whether we should adopt it. A more optimistic take on the relevance of such considerations would consist of taking it that they *explain why* we feel that it is wrong to endorse the consequences of our framework—but the question of whether such hypotheses are true should be addressed experimentally.[64]

Granting such considerations as our objector may, however, she might still be left unhappy with the idea of using our framework to theorize about inquiry. She might still think that she will thereby come to adopt too skeptical a stance on the issue of which questions inquirers have/do not have enough information to settle (not many, that is). For it seems again that our theoretical constructs make it very hard for an inquirer to be in that position. The points from this section are supposed to call the latter impression into question. Their moral is that one does not *have* to adopt a skeptical stance on the target issue in order to embrace our framework.

The first response bears some resemblance to contextualist theories of knowledge-ascriptions and the way they deal with skepticism.[65] But our response doesn't take the semantic contribution of 'knows' to be context-dependent. The

---

[64] In the case of the first consideration, the hypothesis would be that subjects find it wrong to say that an inquirer's information is *not* enough to settle a question *Q* because they take the inquirer's information to be enough to settle a different question, say, the question *p*(*Q*). In the case of the second consideration, the hypothesis would just be an instance of the more general hypothesis that people's judgments about the truth-conditions of natural language sentences are often informed by their judgments concerning the felicity/infelicity of asserting those sentences in context.

[65] See Cohen (1988) and Lewis (1996) for two influential proposals of this sort.

suggestion is rather that the semantic contribution of 'information that is enough to settle *Q*' is context-dependent, much like that of other epistemic modals arguably are.[66]

In particular, the idea is that context will supply a number of presuppositions shared by speakers and hearers that will serve to *restrict* the space of possibilities relative to which *Q* is to be regarded as open or closed by the information possessed by an inquirer.[67] And so there will be situations with the following features: even though the set of possible worlds that constitutes the inquirer's coarse-grained information does not in itself constitute an answer to *Q*—or the inquirer's information doesn't entail any of *Q*'s maximal answers—a proper subset of it *does* constitute such an answer, namely, the proper subset of it containing only those worlds where all our shared presuppositions are satisfied (shared among the interlocutors of our context).

To see how this would work in practice, let us go back to the example featuring Farmer Joe, who was inquiring into *whether it rained last night*. It was part of Joe's evidence, remember, that *the grass is wet*. And we had Joe infer on that basis that *it rained last night*. But then we started to worry about whether the information possessed by Joe—say, as determined by what he knows or what his evidence is—was enough to settle the target question in our favored sense. For it doesn't seem that his information *entails* any of the maximal answers to the target question.

Now, however, with the context-sensitivity of 'information that is enough to settle *Q*' in our hands, we don't *have* to say those things. We don't have to be the skeptics that it seemed we would have to be if we were to adopt the theoretical framework developed here.

For there can surely be contexts where the presuppositions shared by the interlocutors are such as to make a token of the following sentence *literally* true: 'Farmer Joe's information is enough to settle the question of *whether it rained last night*'. There will be no possible worlds compatible with the relevant presuppositions where Joe's grass is wet in the morning, but it didn't rain last night. (Speakers could explicitly express the crucial presupposition here by uttering some sentence like 'If it had not rained during the night, then the grass wouldn't be wet in the morning'). In *none* of those worlds unusual things such as the following happen: someone has gotten out of their way and made a really big effort to trick farmer Joe into thinking that it rained last night by wetting the grass with a hose overnight, a freakish confluence of

---

[66] A canonical semantics for modals in this vein was developed by Kratzer (1977), where the semantic contribution of a modal is also determined by what Kratzer calls the 'conversational background' against which the modal is used (as opposed to the modal itself being *ambiguous*). See Dowell (2017) for more on contextualism about epistemic modals within the Kratzerian tradition.

[67] See Bloome-Tillmann (2014) for a form of epistemic contextualism that explicitly deploys the idea of (pragmatic) presuppositions.

mechanical forces made the water from the river 2km away travel horizontally to Joe's front lawn, etc.

Within that bubble of possible worlds—namely the one that is restricted by the speakers' presuppositions in the envisioned manner—the subset of possible worlds that are compatible with Joe's total body of information *does* constitute an answer to the question of *whether it rained last night*. More formally, where Q is the target question, i_J is Joe's body of coarse-grained information, and p is the set of worlds that are compatible with those contextual presuppositions, $i_J \cap p \in Q$, though $i_J \notin Q$.

Since the sentence 'Farmer Joe's information is enough to settle the question of *whether it rained last night*' is true relative to such contexts, the sentence 'Farmer Joe needs to gather more information first in order to have information that is enough settle the question of *whether it rained last night*' is false relative to those contexts. Full-blown skepticism regarding which questions inquirers have enough information to settle is thereby avoided. Whether a skeptical verdict is to be delivered depends on which presuppositions are shared among the interlocutors who are uttering those sentences.

Switch to a more demanding context, where some of the presuppositions that were relied on in the previous context are lifted, and tokens of 'Farmer Joe's information is enough to settle the question of *whether it rained last night*' become false —but only then, only in those more demanding contexts.

Our notion of information that is enough to settle a question is only as demanding as our context makes it. Strictly speaking, 'information that is enough to settle *Q*' has different contents at different situations of utterance. It is a *context sensitive* expression, in that its semantic value is determined by context (speaker, time of utterance, etc.). Of course, this is a quick and simplified picture of how the semantic values of the relevant expression are to be fixed. In particular, it tells us nothing about how the mechanism of presupposition-lifting works, or about the properties that shared presuppositions need to have in order to restrict the target space of possibilities (the inquirer's body of coarse-grained information). Still, it gives us a good idea of how it would work at the level of semantics.

That is not the only consideration that calls the contention that our framework is over demanding into question. The second one goes as follows. Maybe the objector is simply wrong in thinking that *only rarely* do we have information that is enough to settle a question, even in our favored sense. That of course depends on exactly *how much* information inquirers are typically in possession of. And maybe inquirers are typically in possession of more information than our objector has it. Perhaps the problem is not *our* skepticism about which questions inquirers have enough information to settle, but *the*

objector's skepticism about how much information inquirers are in a possession of in the first place.

In the course of building our models of inquiry, we have mentioned at least two ways in which an inquirer's body of coarse-grained information might be determined: (a) by what the inquirer knows, and (b) by what the inquirer's evidence is.

In the case of (a), the determination goes as follows: take the set *k* of all the propositions *p* such that the inquirer knows that *p*; that inquirer's body of coarse-grained information is then the set of possible worlds where all the propositions from *k* are true. In the case of (b), it goes as follows: take the set *e* of all the propositions *p* such that the inquirer has the evidence that *p*; that inquirer's body of coarse-grained information is then the set of possible worlds where all the propositions from *e* are true.

Of course, under the assumption that a subject's evidence just is what she knows, these two processes (a) and (b) will spit out the same set—but divergence will ensue from the contrary assumption.[68] These were just some examples of how an inquirer's coarse-grained information might be more precisely defined, and still other proposals are possible (the more general idea was, again, that coarse-grained information is determined by fine-grained information).

Now suppose that an inquirer's coarse-grained information is fixed through method (a). Then the more skeptical we already are about knowledge, the more skeptical we will be about which questions inquirers have enough information to settle. If we are not very skeptical about modal knowledge, for example, and we are happy to grant Farmer Joe the knowledge that *if it had not rained last night, then the grass wouldn't be wet in the morning* (read that as a variably strict conditional whose accessibility relation is reflexive), then his total body of information *does* entail that *it rained last night*, seeing as he also knows that *the grass is wet in the morning*.[69] The less knowledge we grant to the inquirer, the more questions there will be such that the inquirer does not have information that is enough to settle them. How skeptical we are about which questions an inquirer has enough information to settle depends on how skeptical we are about knowledge possession.

Similarly, suppose coarse-grained information is fixed through method (b). Then the more skeptical we are about what belongs to a subject's total body of evidence, the more skeptical we will be about which questions inquirer has information that is enough to settle. If only propositions about how things appear to the inquirer count as part of

---

[68] The view that $e = k$ has been influentially advanced and defended by Williamson (2000).

[69] No commitment here to the thesis that subjunctive conditionals in general have the truth-conditions of variably strict conditionals. It just happens that the best sentence we could find through which Joe might express his knowledge in this case is a subjunctive conditional sentence.

his evidence, for example, then there will be fewer such questions (in comparison to the assumption that propositions about how things are in the environment also count as evidence).

And similarly for other ways of determining what an inquirer's body of coarse-grained information is. It is again not the notion of information that is enough to settle a question that *by itself* generates a pessimistic or skeptical view about inquiry. For it is always possible to adjust our total theory of inquiry elsewhere to make it more common-sense friendly, without even changing its notion of information that is enough to settle a question.

§5.4 Where we counter those objections to the Go Gather Norm

We saw initially that there was a serious worry about the Go Gather Norm (GG) when its notions are interpreted in the way we did, namely, that it is an overly demanding norm, on point of it being too hard for us to possess information that is enough to settle the questions we want to settle.

Furthermore, it seems that we often wrap up inquiry by inferring one of the question's answers through abductive reasoning from our evidence—and quite rationally so—*despite* not having information that is enough to settle the question in our sense. The considerations from the two previous sections now provide us with a variety of responses to these objections. The norm under discussion here is, to repeat:

(GG) If one wants to settle *Q*, but one's coarse-grained information is not yet enough to settle *Q*, then one should see to it that one gathers new information that bears on *Q*.

Now when an inquirer seems to abductively infer an answer to a question *Q* that is made probable by her information, though it is not yet settled by it in our sense, we can say that the inquirer has thereby settled yet another question, and one that she did have enough information to settle, such as the question *p*(*Q*) of *which of Q's answers are more probable than the others*. For, presumably, when she makes one such inference, she believes of the answer that she inferred from her evidence that it is more probable than the other answers, implicit as such a belief may be. So there is something that the inquirer did right to the eyes of (GG), even under our interpretation of 'information that is enough to settle a question'. Of course, there is also something that she did wrong, namely, she did not gather more information that bears on *Q*. That is still something that she should have done, at least assuming that she really didn't have information that was enough to settle *Q* and that she wanted to settle *Q*. If this is what she wanted, then that is what she should have done.

But the desired effect of the present observation is that our inquirer did not do *everything* wrong. She has violated (GG) when it comes to *Q*, but she did not violate (GG) when it comes to *p*(*Q*) (she had information that was enough to settle the latter question). It is one right and one wrong, as opposed to just one wrong.

Furthermore, in many such cases it is *felicitous* to say that the inquirer has information that is enough to settle *Q*—say, because he has information that makes it rational for him to accept one of the maximal answers to *Q* and act as if it is true for some practical purposes. The felicity of the utterance gives us then the impression that the inquirer does *not* have to gather more information, in contrast to what (GG) says. But those are just the consequences of what is felicitous for us to assert, not of what is true for us to assert.

Those are the responses borrowed from the considerations from §2.2. Based on what we saw in §2.3, however, we now say that it might not only be *felicitous*, but also *literally true*, to say that the inquirer has information that is enough to settle *Q* in the relevant cases—in which case the antecedent of the conditional embedded in (GG) is not even true, and therefore the inquirer is not violating that norm with respect to *Q*.

We have seen two sources of that possibility. First, whether a sentence like 'The inquirer has information that is enough to settle *Q*' is true depends on the presuppositions that are shared among the interlocutors at the context at which that sentence is uttered. If those presuppositions do enough to restrict the space of possibilities to be quantified over, then the inquirer will have information that is enough to settle *Q*, and (GG) (as uttered in the same context) will not issue any reproach to that inquirer, even if the inquirer does not collect new information that bears on *Q*.

Second, how many questions are such that the inquirer has enough information to settle depends, among other things, on how much the inquirer knows or how much evidence he possesses. The more he knows, or the more evidence he has, the more questions will he have enough information to settle, *ergo* the less questions will there be such that he violates (GG) with respect those questions. For a skeptic about how much we know/how much evidence we have, (GG) will be issuing reproaches all over the place. But it doesn't have to be that way. Instead of demanding the notion of *information that is enough to settle a question* that is at play in (GG) to be more inclusive, then, we can rather demand (*contra* the skeptic) that the notions of *knowledge* and *evidence* be more inclusive

In any case, these are some of the important moving parts of our truth-conditions for 'the inquirer has information that is enough to settle *Q*'—namely, contextual parameters and how much information inquirers typically possess—and the whole theory of inquiry of which (GG) is part only gets to have a more determinate face once those moving parts are fixed in one way or another. But, in the end, it is perhaps

not very problematic to grant that inquirers often violate (GG) (because they often abide by (GG) regarding other questions, because it is often felicitous to assert that their information is enough to settle questions), and *negotiable* whether they indeed violate it that often (depending on how much is presupposed in the context of utterance, depending on how much inquirers know). We have seen no fatal objections to our norm and the theoretical framework that backs it up here, which includes our explication of the notion of *information that is enough to settle a question*.

Before we move on to the next chapter, however, let us now briefly look at a slightly different kind of worry about (GG). What if it is very difficult or even impossible for the inquirer to acquire information that is enough to settle *Q*, on account of the very absence of such information at the time/world the inquirer is in?

Suppose a biographer wants to settle the question of *what Marie Skłodowska Curie's favorite dish was*, but there are no entries on this topic in Marie's notebooks and the letters she sent to other people, etc.—no records whatsoever that our biographer could consult through which he could come to learn what Marie's favorite dish was (it is a gap in history that won't ever be filled). Suppose furthermore that our biographer's information is not enough to settle that question. So the antecedent of the conditional embedded in (GG) is true of him. Yet, it doesn't seem that he should go looking for information that will settle the issue, simply because *there isn't* such information for him to acquire. Why should one try to do the *impossible?* One shouldn't. But (GG) says one should, so (GG) must be wrong.

But the natural response is that, in cases such as this one, what the inquirer should do is just stop wanting to settle his question. In general, agents should not want to do what is impossible for them to do. In particular, inquirers should not want to settle *Q* when it is impossible for them to settle *Q* (see also the Anti-Impossibility Norm of Inquiry in §8.2). And not wanting to settle *Q* is indeed a way of coming to abide by the norm (GG). Even though that norm says that one should go gather information when certain conditions are satisfied, it also gives one the option of not satisfying those conditions to begin with. In the case of the biographer we just saw, he should not want to settle the question of *what Marie Skłodowska Curie's favorite dish was*. (And, of course, if he wants things that he shouldn't want, he will be bound to do things that he

shouldn't do by the true norms. But those are just consequences of violating the norms to begin with).[70]

So the case doesn't threaten the correctness of (GG) after all. For the 'should' operator of that norm, remember, is to be read as having wide-scope. Its semantic contribution ranges over the whole conditional: *if one wants to settle Q and one's coarse-grained information is not yet enough to settle Q, then one gathers new information that bears on Q*. Assuming as we are that the target conditional is a material conditional, what (GG) says is that one should be such that: *either it is not the case that (one wants to settle Q and one's coarse-grained information is not yet enough to settle Q), or one gathers new information that bears on Q*. So there is at least two ways in which one can abide by (GG) with respect to a question *Q*: either by *not* wanting to settle *Q*, or by gathering new information/making it so that one's coarse-grained information is enough to settle *Q*.

The biographer from the case described in the previous paragraph doesn't *have* to go gather more information about Marie Skłodowska Curie in order to bring himself to abide by (GG)—he can instead cease to want to settle the question of *what Maria Skłodowska Curie favorite dish was*. And, assuming again that it is impossible for him to settle that question, this seems indeed to be the thing for him to do. Such a verdict is perfectly consistent with (GG).

We have submitted the Go Gather Norm (GG) to criticism, and we have found it to withstand scrutiny. Inquirers who want to settle a given question, though they do not yet have information that is enough to settle that question, should gather more information. We now turn to a critical assessment of the Go Figure Norm (GF).

---

[70] The situation wouldn't change much if *there were* still a bit of information regarding Marie Curie's favorite dish for one to gather—not if we continue to assume that it is still impossible for one to finally settle the question. Sure enough, if there is still some such bit of information, then one can still make some incremental change, say, in one's credences regarding *what Marie Curie's favorite dish was*—but then there will accordingly be questions about what the probability is that this or that was Marie Curie's favorite dish such that it is *not* impossible for one to know the answers to those questions.

Chapter 6

§6.1 Where we also submit the Go Figure Norm to critical scrutiny

In §4.4 we have introduced two norms of inquiry, the Go Gather Norm (GG) and the Go Figure Norm (GF). We noted that, given that an inquirer wants to settle *Q*, one or the other of these norms will require him to do something—either go gather more information or rather process the information he already has. (GG) tells him to do the former if his information is not yet enough to settle *Q*, and (GF) tells him to do the latter if his information is already enough to settle *Q*.

The Go Figure Norm (GF) is more specific about the kind of action (in a broad sense) the inquirer is supposed to perform. Here is that norm again:

(GF) If one wants to settle *Q*, and one's coarse-grained information *i* is enough to settle *Q*, then one should see to it that one competently deduces some fine-grained information *x* such that $v(x) = a$, where *a* is the maximal complete answer to *Q* such that $i \subseteq a$.

Suppose you satisfy the antecedent of the conditional embedded in (GF) with respect to some question *Q*, say, the question of whether *you could have had different parents (while still being you)*. So your coarse-grained information constitutes one of the complete answers to that question—presumably unbeknownst to you, seeing as you want to know what the answer to that question is. That complete answer is a subset of one of the *maximal* answers *a* to the target question (in this case, the *no* answer). So (GF) tells you to see to it that you deduce some fine-grained information that carries that information *a*, so that you thereby acquire a maximal *fine-grained* answer to your question, over and above the complete coarse-grained answer to it that you already have. That would involve adding a vehicle *x* that carries the coarse-grained information that *a* into your representation of reality, that is, a vehicle *x* such that $v(x) = a$.

(A maximal fine-grained answer to *Q* is a bit of fine-grained information, that is, a pair of a vehicle and the content it carries. It is connected to *Q* in virtue of the second element as follows: the content it carries is a maximal answer to *Q*).

We initiate our assessment of (GF) by getting the following objection out of our way, on account of it being similar to the last objection to (GG) that we discussed in the previous chapter. What if one wants to settle *Q*, and one's information is enough to

settle *Q*, but one has *already* settled *Q* by deducing a fine-grained answer to it? In cases of this sort, it doesn't seem as if one should again deduce the fine-grained answer to *Q* that one is already in possession of. This form of 'double concluding' seems otiose.

Here is another way to put the objection. An inquirer will have some representation or other of how the world is like—her internal model of reality. There are gaps in that representation, and inquiry is supposed to fill some of those gaps. Where the inquirer has already competently added representational vehicles to her total representation of reality, however, she has already filled some of the gaps in it. And then it makes no sense for her 'to add' those very representational vehicles to her representation of reality yet again. For there is no more gap to be filled there in the first place, or nothing to add that isn't already there.

The response to this objection should be more or less obvious by now. (GF), like (GG), is a wide-scope norm. The semantic contribution of 'should' in it ranges over the whole (material) conditional *if one wants to settle Q, and one's coarse-grained information i is enough to settle Q, then one sees to it that one competently deduces some fine-grained information x such that v(x) = a* (*where a is the answer to Q with* $i \subseteq a$). In order to bring herself to abide by (GF), then, the inquirer doesn't *have* to deduce a fine-grained answer to her question from the information available to her. She can instead stop wanting to settle her question. And, arguably, if she has already competently deduced a fine-grained answer to *Q*, she shouldn't want to settle *Q* anymore (at least assuming that the target fine-grained answer is an answer to *Q as presented under the same guise* under which one wants to know *Q*—see §6.4 for this important detail that is now left implicit).

The problem here is not that the inquirer wants *the impossible*, as in the case from §5.4, but rather that she wants *what she already has*. Just like before, however, it seems that what the inquirer should do all things considered is not want to settle *Q* anymore, and this verdict is perfectly consistent with (GF).

A different kind of worry about (GF) stems from considerations about cognitive limitations. What if the inquirer wants to settle *Q*, and her information is enough to settle *Q*, but she lacks the skills that she needs to competently deduce a fine-grained answer to *Q*? Under these circumstances, it seems that the inquirer should not try to competently deduce such an answer from the information possessed by her, simply because she won't manage to do that if she tries to. If she were to try to deduce such an answer, she would just waste her time and cognitive resources, attempting to do something that she is not really in a position to do. That would be counterproductive,

not at all instrumentally rational—and yet we have postulated that the 'should' of (GF) and (GG) is an instrumental 'should'.

The challenge is welcome, for it invites clarification about what kinds of activity count as *seeing to it that one competently deduces* a certain bit of fine-grained information. In addition to the activity of actually deducing the relevant bit of fine-grained information, there are also activities that allow the inquirer to acquire new reasoning skills and thereby come to be in the position to perform deduction.

Suppose the inquirer lacks the skills that would allow her to competently deduce a bit of fine-grained information *x* of a certain kind from the information possessed by her. Let us say that, in principle, it would be possible for *someone* to deduce *x* on the basis of the total body of information possessed by the inquirer—but the inquirer herself is not able to do that. She can then see to it that she competently deduces *x* by developing new inferential skills that *will* eventually allow her to competently deduce *x*. If she does that, she abides by (GF), even though she does not yet deduce *x* from the information possessed by her, and she still wants to settle *Q*. (The acquisition of such new skills could involve, for example, the development of pattern recognition of relations among symbols/strings of symbols of a certain type).[71]

There is more than one way in which one can see to it that one competently deduces a fine-grained answer to a question, then: actually deducing it, on the one hand, and learning how to deduce it, on the other. The case of an unskilled reasoner constitutes no objection to (GF) after all. There is something that she can do to abide by that norm, too, even holding fixed the fact that she wants to settle her question (a question she already has enough information to settle). Of course, even *after* the inquirer acquires the skills to competently perform the relevant deduction, her job is not completely done yet, in that she should still see to it that she performs the deduction. But the point is that it wasn't the case *before* she acquired the relevant skills that the norm was telling her to try to competently deduce the answer to her question, which would indeed be a waste of her time and cognitive resources, seeing as she was not in a position to competently perform the deduction.

The norm makes perfect sense, again as a matter of instrumental rationality. Suppose the inquirer wants to settle *Q*, and she has information that is already enough to settle *Q*. Unfortunately, however, she doesn't have the ability to deduce the fine-grained answer to *Q* that can be read off of the coarse-grained information possessed by her. Given what she wants, the thing for her to do seems to be to acquire that ability—for that will put her in a position to actually settle the question she wants to settle. That is indeed a means to her end (much like gathering more information is a means to

---

[71] See Landy, Allen and Zednik (2014) for a perceptual account of symbolic reasoning.

her end when she wants to settle a question without having information that is enough to settle it yet).

§6.2 That the 'should' of our norms of inquiry doesn't mean the same as 'must'

Still, we have to be careful about how to interpret the instrumental 'should' of our norms of inquiry (GF) and (GG), this time regarding its *strength* rather than its scope (which we already set to be wide).

We can find the desired strength to assign to 'should' by contrasting it with 'must'. Where 'must' expresses strict obligation, we have good reasons to reject the following norm:

(GF+) If one wants to settle *Q*, and one's coarse-grained information *i* is enough to settle *Q*, then one must see to it that one competently deduces some fine-grained information *x* such that $v(x) = a$, where *a* is the maximal complete answer to *Q* such that $i \subseteq a$.

The only difference between (GF) and (GF+) is that the former deploys 'should' and the latter deploys 'must'. (GF+) seems to be false, on account of the following kind of consideration. Suppose the inquirer wants to settle *Q* and he already has information that is enough to settle *Q*. He does have the skills that would allow him to settle *Q* via deduction on the basis of information he already has—but it would be incredibly difficult for him to do that. Why *must* he then see to it that he performs the target deduction himself (or rather not want to settle *Q* anymore), instead of, say, consulting his epistemic superiors on the topic?

For example, *Q* could be the question of whether a certain formula is a theorem of (some version of) set theory such that, even though our inquirer does have the skills to competently decide whether it is a theorem or not, it would be very difficult and time-consuming for him to do that. But his partner, who happens to be by his side, is a set-theorist—so that he could just ask *her* whether the target formula is a theorem of set theory and call it a day (she knows the answer). Again, seeing as our inquirer can easily obtain the answer to his question by consulting an epistemic superior, it seems that it is not the case that he *must* see to it that he deduces the answer to his question himself, even though it is possible for him to do that.

The moral is that, if (GF) is to be true, the word 'should' in it must not make the same contribution that 'must' does in (GF+). For (GF+) is false. More generally, the 'should' of our norms of inquiry better not express any strict kind of obligation. Inquirers need not be under any kind of obligation to satisfy the conditionals embedded in them.

A contrast is to be drawn, then, between ‘should’ and ‘must’ as they occur in our norms of inquiry. The former is to be *weaker* in strength than the former, in that ‘should(*p*)’ does not entail ‘must(*p*)’, even though ‘must(*p*)’ entails ‘should(*p*)’ (the special case of our focus is one where ‘*p*’ is a material conditional sentence).

We draw the relevant contrast semantically, and we do it while remaining within the framework of possible world semantics. At bottom, the solution is that ‘should’ quantifies over a *proper subset* of the set of possible worlds that ‘must’ quantifies over. They both quantify over possible worlds containing ideal inquirers—but the inquirers that belong to the worlds that ‘should’ quantifies over are more idealized than the inquirers that belong to some of the worlds that ‘must’ but not ‘should’ quantifies over.

‘Must(*p*)’ is true just in case ‘*p*’ is true at all worlds in *M*, whereas ‘should(*p*)’ is true just in case ‘*p*’ is true at all worlds in *S*, where $S \subset M$. At all worlds of *M*, inquirers always do what they must to achieve the goals of inquiry. At all worlds of *S*, they always do what they should to achieve them. We can represent the relation between these two spaces of possibility diagrammatically thus:

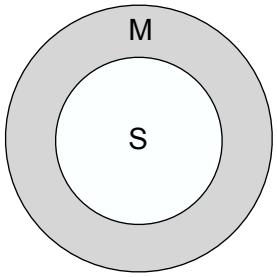

The inquirers that belong to the worlds in *S* (white area) are very different from us human inquirers. They are ideal along the dimension of instrumental rationality, at least when it comes to the activity of inquiry. They are more diligent in their search for truth, and they don’t have the cognitive limitations that we have. In particular, and relevant to our present concerns, they promptly deduce a fine-grained answer to any question they want to settle, given that their information is already enough to settle that question. If it isn’t, they promptly set themselves off in the search for more information —which includes asking other inquirers for information—or rather they cease to want to settle their questions. But they only consult others when they don’t already have information that is enough to settle the questions the want to settle.

The worlds in *M* that are not in *S* (grey area) feature inquirers that are also idealized—though they are not as supremely idealized as the inquirers from the possible worlds within *S*. The former inquirers have more cognitive limitations than the

former ones. One of the consequences of this is that the more resource-bounded inquirers from the grey area within *M* will sometimes resort to epistemic superiors in their attempt to settle their questions, even though they have information that is enough to settle those questions. For, given their constraints, it would still be taxing and time-consuming for them to settle the target questions on their own, using their own reasoning skills. The denizens of the *S* area don't have any need for that, however. Both the inquirers from *S* and the ones from *M* always get their inquiry-related things done in the right way, and in that they are ideals for us to approximate, though they sometimes do this in different ways.

With such a semantics in place, it becomes *harder* for 'must(*p*)' to be true than it is for 'should(*p*)' to be true. The truth of 'should(*p*)' is cheaper than the truth of 'must(*p*)'. Accordingly, 'should' has less normative force than 'must': if even the *not so supremely ideal* inquirers satisfy *p*, then inquirers like us (human beings) are more under pressure to make *p* true (a case where we *must* make *p* true), in comparison to a case where only supremely ideal inquirers satisfy *p* (a case where we *should*, though it is not the case that we *must* make *p* true). There is less pressure for us to approximate the inquiry-related procedures of the supremely ideal inquirers of the *S*-worlds than there is for us to approximate the inquiry-related procedures of the *less than supremely ideal* inquirers of the *M*-worlds that are not *S*-worlds.

We don't need to work out all the details of such a semantics for 'must' and 'should' as they occur in our norms of inquiry. The important thing to note now here is that it allows us to circumvent the kind of objection that concerned us above. It does that by validating (GF) without validating (GF+). Even though the supremely ideal inquirers from *S* always go on and deduce the fine-grained answers to the questions they want to settle (assuming that their information is enough to settle them), again, some of the less than supremely ideal inquirers from *M* don't always do that. As a result, the conditional embedded in (GF) is true at all members of *S*, but false at some of the members of *M*. Since (GF+) embeds that very same conditional under 'must', it comes out false, even though (GF) comes out true.

Most importantly, the fact that an inquirer can just resort to an epistemic superior, instead of making the deduction by herself, doesn't threaten the truth of (GF). For there might be less than supremely ideal inquirers in some of the worlds in *M* that do not make the deduction by themselves either—they consult their epistemic superiors instead. Even though it is still the case that the inquirer *should* make the deduction (given the assumption that she wants to settle a question she already has enough information to settle), what she *must* do is *either* make the deduction *or* consult an epistemic superior. That is, she would have to make the deduction in order to be like

the supremely ideal inquirers from *S*, but she doesn't have to make the deduction in order to be like the less than supremely ideal inquirers from *M*.

There should still be other candidate semantics that do the job of telling 'should' and 'must' apart in the envisioned manner, in such a way that the fact that inquirers can settle their questions by consulting their cognitive superiors (in case there are any), rather than by deducing the answer by themselves, constitutes an objection only to (GF+) but not to (GF). We saw but just one way of building such a semantics by way of illustration. We will get back to what makes our norms featuring 'should' true in §§9.1–9.2, where we ground those norms on facts about the detraction and accretion of instrumental value to inquiry, relative to its very constitutive goal.

If the reader still thinks that (GF) should be discarded, or that there shouldn't be a reading of 'should' that makes it true to begin with, we can add conditions to the antecedent of its conditional such as '… and there is no epistemic superior for the inquirer to consult about whether *Q*', or rather disjunctions to its consequent such as '… or the inquirer consults an epistemic superior about whether *Q*' and thereby get a new norm that the reader will hopefully agree with. But we will run with (GF) as is—the morals to be drawn here wouldn't change dramatically if we beef its antecedent or its consequent up in this way.

§6.3 Where we start integrating different norms of inquiry into a unified system of norms

We can now start thinking about how the Go Gather Norm (GG) and the Go Figure Norm (GF) fit into a more comprehensive system of norms of inquiry. At a very general level, those two norms are again exhaustive vis-a-vis what an inquirer is to do, depending on the situation she finds herself in.

Given that the inquirer wants to settle some question, what she is supposed to do depends on how her total body of information bears on that question. If the question is still open relative to her total body of information, (GG) tells her to gather more information. If the question is already closed relative to her total body of information, (GF) tells her to flesh out the answer from that body of information—to make explicit what was only implicit (perhaps the philosopher is often in this kind of situation *qua* inquirer). The first norm directs the inquirer to perform information-gathering inquiry, whereas the second one directs her to perform armchair inquiry, though again it doesn't follow that the inquirer overall *must* or has a duty to do those things.

So far so good—but how do those norms relate to norms that tell the inquirer *not* to inquire into questions? Do they form a consistent set of norms? And can the constructs that we have deployed to formulate and justify (GG) and (GF) also be used to justify such prohibitive norms? Such are the issues that will occupy us from now on.

In §3.1 we briefly discussed a prohibitive norm of inquiry that was vindicated by the possible worlds model of inquiry, namely:

(N) One shouldn't gather more information in an attempt to settle *Q* if one already knows *Q*.

This norm continues to be justified in the same way by the new model of inquiry, which also deploys the notions of coarse-grained information and questions from our framework (though it features additionally constructs for fine-grained information and reasoning skills). If the inquirer already knows *Q*, then his coarse-grained information is *already* one of *Q*'s answers, and no further chopping of that body of coarse-grained information—a space of possibilities that are open to the inquirer—will improve the situation of how his information bears on *that* question. The point of gathering new information just is that of chopping down that space of possibilities, however. So collecting more information in an attempt to settle *Q* would be a waste of the time and resources of that inquirer, *ergo* (N).

Norm (N) is strictly concerned with information-gathering inquiry—but we can generalize it to any kind of inquiry, including armchair inquiry. Its generalization would be the Ignorance Norm of Inquiry:[72]

(INI) One shouldn't inquire into *Q* if one already knows *Q*.

Since the proposition that *one gathers more information in an attempt to settle Q* entails the proposition that *one inquires into whether Q*, (INI) entails (N). But since the proposition that *one inquires into whether Q* does not entail the proposition that *one gathers more information in an attempt to settle Q*, (N) does not entail (INI). We want our system to allow us to derive more specific norms of inquiry from a few more general ones—so (INI) is our bigger fish here.

---

[72] The norm is named this way here after Whitcomb (2017), who fleshes out an analogous norm for the speech act of *asking*, and Friedman (2017), who fleshes out an analogous norm for *interrogative attitudes*. Still concerning the speech act of asking, Hawthorne (2004, p. 24) writes: 'At least in the normal case, we shouldn't ask a question if we already know the answer'. Whitcomb (2017) argues that inquisitive asking is constitutively governed by an ignorance norm. Since our use of the verb 'inquire' doesn't denote a speech act of asking, Whitcomb's norm and (INI) are not quite the same. And since we are understanding inquiry as an *activity* here (and a goal-directed one at that), our (INI) isn't equivalent to Friedman's (2017) ignorance norm, either. For, as we saw in §1.2, inquiry in our sense is not an attitude or even a combination or profile of attitudes. And the norm presented in Friedman (2017) puts constraints on *intensional attitudes* of a certain sort, namely interrogative attitudes such as that of being curious or wondering.

Semantically, we assign wide-scope to ‘should’ here, just like we did before. So (INI) has the logical form $S\neg(IQ \wedge KQ)$, where ‘S’ is the should-operator, ‘$\wedge$’ is the sign for conjunction, ‘IQ’ translates ‘one inquirers into Q’ and ‘KQ’ translates ‘one knows Q’. What (INI) says, then, is that the inquirer should not be in a situation where she both inquires into Q and knows Q at the same time and possible world (parameters left implicit in our formulation). Equivalently, with the material conditional: $S(IQ \supset \neg KQ)$. That is, one should be such that one inquires into a question Q at a certain time and possible world only if one does *not* already know Q at that time and world. One has to *ignore* Q (ignore *whether such-and-such is the case*, ignore *who did this-and-that*, etc.), whence the name ‘Ignorance Norm of Inquiry’.

How can we justify (INI) using the constructs that we have been deploying so far to theorize about inquiry? Suppose again that the inquirer knows Q. So there is an answer a to Q such that a is true at all possible worlds that constitute that inquirer’s coarse-grained information. It follows from our definitions that the inquirer’s information is enough to settle Q—so there isn’t any need for him to gather more information in an attempt to settle Q. Furthermore, the inquirer represents the world he is in as being such that a, which means that there is some representational vehicle that constitutes a fine-grained answer to Q that is part of the inquirer’s representation of reality (some bit of fine-grained information that carries the coarse-grained information that a). So there is no need for him to engage in armchair inquiry into Q, either—for he is already in possession of a fine-grained answer to Q.

Under the assumption that the inquirer knows Q, then, engaging in either of information-gathering inquiry or armchair inquiry into Q would be a misuse of that inquirer’s time and resources. Those are counterproductive activities for our inquirer to perform (he could be investing his time and resources to settle other questions he wants to settle). Therefore, an inquirer shouldn’t at the same time and world know Q and inquire into Q, again in the instrumental sense of ‘should’. Inquiring into Q when one already knows Q is like trying to go to Paris when one is already there.[73]

For example, most likely you already know *which city you were born in*. That means that you already know that *you were born in x*, where *x* is the city where you were born. (INI) then says that, in this scenario (where you know the answer to that question), you are doing something that you shouldn’t be doing if you also inquire into *which city you were born*. That activity is instrumentally counterproductive to you.

---

[73] Related to the point of the previous footnote, note here that this justification of (INI) is made on different grounds from those on the basis of which Friedman (2017) justifies her ignorance norm for interrogative attitudes—namely, that an inquirer who violates that norm thereby holds mutually incoherent attitudes.

(INI) and its accompanying justification have their wrinkles, however, and not small ones at that. We turn to these now.

§6.4 That the Ignorance Norm of Inquiry, as presented above and justified in the way it was, leaves a number of important details implicit

First, the justification we have provided above assumes that, whenever one knows that *p,* one is thereby in possession of fine-grained information that carries the (coarse-grained) information that *p*.[74] For it concludes that *a representational vehicle that constitutes a fine-grained answer to Q is part of the inquirer's representation of reality* under the supposition that *the inquirer knows Q*.

But what about implicit knowledge? Presumably, where ascriptions of implicit knowledge are called for, the knower doesn't necessarily have a representation of the fact that he is said to know in his internal model of reality. What is known in such cases is only implicit in his internal model of reality. And it seems that questions whose answers are known only implicitly deserve inquiry, too, so that this knowledge can be brought to the surface or made explicit in the inquirer's thought, thus giving him the opportunity to share this knowledge with others through language and creating new inferential opportunities both to him and his interlocutors (inference as vehicle-manipulation). It doesn't seem instrumentally irrational, for example, for linguists to make their implicit knowledge of which constructions are grammatical explicit.

The quickest fix to this first concern is to think of the ascription of knowledge in (INI) as an ascription of *explicit* knowledge, or knowledge that provides the inquirer with a representation that constitutes a fine-grained answer to the target question. And similarly for the ascription of knowledge that starts off the *justification* we have provided for (INI) in the previous section (the justification started with 'Suppose again that the inquirer knows *Q*…').

In this way, our justification of (INI) ends up relying on some of the same considerations that we have used to dismiss the first objection to (GF) from above (§6.1), which had us consider the case of an inquirer who wants to settle *Q* though she has already settled *Q*. If the inquirer has *already* closed a gap in her model of reality, why want to know how to fill that gap yet again? That would be an otiose want. In such cases, we said, one shouldn't want to know *Q*.

And here, similarly, if the inquirer has already closed a gap in her model of reality, and she did so by gaining explicit knowledge of a question's answer, then why

---

[74] A bit of fine-grained information is, remember, a pair <*x*, *v*(*x*)> of a vehicle *x* and the coarse-grained information *v*(*x*) carried by that vehicle (some set of possible worlds). We say that a bit of fine-grained information carries some coarse-grained information, albeit a bit redundantly, when its constitutive vehicle carries that coarse-grained information.

inquire into that question again? That would be an otiose inquiry, for explicit knowledge closes the gap, settles the question. One shouldn't inquire into *Q* under those circumstances, and that is just what (INI) says. (INI) does not forbid inquiry into questions whose answers are known only implicitly. The linguist, the mathematician, the philosopher—to the extent that they are inquiring into questions whose answers they already know, although in an implicitly manner—are not thereby doing something wrong to the eyes of (INI).

Second, what if the inquirer already knows *Q* under one mode of presentation or guise, and he goes on to inquire into *Q* under a new guise? (We are thinking of the guises as the representational vehicles themselves). That need not be a counterproductive form of inquiry.

For example, there is just *one* question whose maximal complete answers are respectively the *total* set of possible worlds (the necessary truth) and the empty set (the necessary falsehood), though it may present itself to the inquirer under different guises.[75] One can approach the same question under different guises, and know the answer to it as presented under one guise without knowing the answer to it as presented under an alternative guise.

Settling the same question under a new guise might again provide the inquirer with cognitive benefits that were not afforded by her settling it under a previous guise. In the realm of necessary truths—so our framework for theorizing about inquiry has it—what really makes a cognitive difference for us are the representational vehicles that have those truths for their coarse-grained information. That is, what makes a cognitive difference are the guises under which those truths are presented (what is *new* about the new model of inquiry is just the function *r* that reads off vehicles of the inquirer's coarse-grained information).

Suppose the inquirer explicitly knows the necessary truth that *t* under a guise *g*, which is a representational vehicle that carries the coarse-grained information that *t*. Say that *g* is (an internalized token of) the linguistic representation '37 is a prime number'. The inquirer knows that *t* under *that* guise. But he has been thinking a lot about numbers lately, and now he starts inquiring into that same question of *whether t* but under a different guise, say, 'Is π a rational number?'.

Now since our inquirer is not yet in possession of a fine-grained answer to that question *as presented under the latter guise*, it is not instrumentally counterproductive

---

[75] Such questions would be represented as follows in the framework we have adopted—questions as the downward closed sets of their maximal (coarse-grained) answers: {$\varnothing$, *W*, …}, where $\varnothing$ is the empty set, *W* is the set of all possible worlds, and the ellipsis stands for all proper subsets of *W* and the subsets of all those sets, all the way down to the singletons {*w*}, for any $w \in W$.

for him to inquire into it under that guise—even though he knows the answer to that question as presented under a different guise. He has a bit of fine-grained information that constitutes a fine-grained answer to the question of *whether t* as expressed by 'Is 37 a prime number?', but he doesn't have a bit of fine-grained information that constitutes a fine-grained answer to that question as as expressed by 'Is π a rational number?'.[76]

The question of *whether 37 is a prime number* is the same as the question of *whether π is a rational number*, though that same question can present itself to an inquirer under different guises. Our inquirer knows *whether 37 is a prime number* when that question is presented under the guise 'Is 37 a prime number?', but he doesn't know *whether 37 is a prime number* when it is presented under the guise 'Is π a rational number?'.

So we seem to have a case in front of us where the inquirer knows *Q* and inquires into *Q* at the same time, but he is not doing anything wrong or wasteful of his time and resources. After all, acquiring a fine-grained answer to the question of *whether t* as presented under the new guise (say, 'Is π a rational number?') may give him new inferential opportunities that were not afforded by his knowledge of the fine-grained answer to that same question as presented under the old guise (say, 'Is 37 a prime number?'). It may be, for example, that the former but not the latter will allow the inquirer to promptly infer the fine-grained answer to the question presented under the guise 'Does the decimal expansion of π eventually terminate?'. That benefit is afforded by the inquirer's knowledge of the answer to her question under one guise, not afforded by her knowledge of the answer to that question under the other guise.

Summing it up, it is not pointless or counterproductive for the inquirer to inquire into a question *Q* such that he already knows *Q*. So (INI) is not a true norm of inquiry after all—it is not the case that inquirers should abide by it, in the instrumental sense of 'should'.

The solution, however, and the detail that is missing from our formulation of (INI) from above, is more or less obvious. (INI) should actually be read as saying that one shouldn't inquire into *Q*, *as presented under a certain guise g* when one already

---

[76] Defenders of the Russellian view of propositions, for example Salmon (1995), frequently deploy some such notion of *guise* or *mode of presentation* to shield their view against certain objections and make sense of cognitive/epistemic phenomena. It is not only Russellians who posit guises, of course. That theoretical construct is obviously useful also to those who make use of (sets of) possible worlds to theorize about mental content and assign them to be the semantic contribution of 'that'-clauses in ascriptions of intensional attitudes. See Williamson (2020, pp. 246–250) and Kratzer (2022) for two examples of that from the more recent literature.

knows *Q as presented under that same guise g*.[77] But there need not be anything instrumentally off, at least not as far as (INI) itself goes, with an inquirer who inquires into *Q as presented under guise g* while knowing *Q as presented under some different guise h*.

A guise is a representational vehicle with a particular format, so that different guises can carry the same coarse-grained content and still not be the same. The same question can be carried by different vehicles, or be presented under different guises. Just like there are fine-grained *answers* to questions (pairs of vehicles and coarse-grained answers), then, so there are fine-grained *questions* (pairs of vehicles and coarse-grained questions).

The inquirer can be sensitive to differences between two interrogative sentences without being sensitive to their commonality of content (the two interrogatives carry the same coarse-grained question, understood here as a set of coarse-grained answers). When that is the case, the two interrogative sentences constitute two different guises through which the same question presents itself to the inquirer. Since the inquirer doesn't realize that she is dealing with the same question under two different guises, however, it won't be instrumentally counterproductive for her to inquire into it under one guise while knowing the answer to it under a different guise. But, when properly understood, (INI) does not say otherwise.

This latter fix on (INI) also shows us that (GF) hides a guise-parameter, too. We are to read (GF) as telling the inquirer to see to it that she deduces a representational vehicle *x* that constitutes a fine-grained answer to *Q as presented under guise g*, given that he has enough information to settle *Q* and he wants to settle *Q as presented under that same guise g* (the representational format through which the inquirer entertains that question and desires to settle it).

So the guise under which the inquirer wants to know the answer to the question must be paired with the guise under which he is supposed to deduce that answer. We reformulate our norm and make these parameters explicit as follows, then:

(GF) If one wants to settle *Q as presented under guise g*, and one's coarse-grained information *i* is already enough to settle *Q*, then one should see to it that one

---

[77] And similarly, where we wrote before in §6.1 about (GF) that '… arguably, if she [the inquirer] has already deduced the target fine-grained answer to *Q*, she shouldn't want to know *Q* anymore', we are to interpret that as saying that the inquirer shouldn't want to know *Q as presented under guise g* if she has already deduced a fine-grained answer to *Q as presented under that guise g*. The fix wasn't strictly needed there, because the objection that was being raised against (GF) already assumed that some such violation (of having deduced a fine-grained answer to *Q* under *g* while wanting to know *Q* under *g*) was already in place—and all we did was to show that (GF) was consistent with such a situation and our other normative judgments about it.

competently deduces some fine-grained information *x* such that *v*(*x*) = *a*, where *a* is the maximal complete answer to *Q* such that *i* ⊆ *a*, and *x* constitutes a fine-grained answer to *Q as presented under guise g*.

What is it for a representational vehicle *x* to constitute a fine-grained answer to *Q* as presented under *g*? It must at least be the case that *v*(*x*) = *a* for some maximal member *a* ∈ *Q*. In addition to that, *x* must be what we might call a *declarative rendering* of *g*. For example, 'The universe had a beginning' and 'The universe did not have a beginning' are declarative renderings of the interrogative 'Did the universe have a beginning?'. Or we might think of a diagram displaying a circle of possibilities one of the sub-areas of which is scratched (standing for the *elimination* of those possibilities) as a declarative rendering of an interrogative diagram that simply divides that circle into different sub-areas, without scratching any of them (representing the question of whether the actual world belongs to one or the other sub-area of possibilities).

§6.5 Where we also make those details of the Ignorance Norm of Inquiry explicit and we introduce yet another purported norm, namely, the Knowledge Norm of Inquiry

We can now trim the Ignorance Norm of Inquiry (INI) a bit further. A more precise formulation of that norm, which explicitly takes the latest details into account, goes as follows:

(INI) One shouldn't inquire into *Q*, as presented under guise *g*, if one explicitly knows *Q* as presented under *g*.

That is more precisely what we are thinking of when we endorse the Ignorance Norm of Inquiry.[78] Formulations of (INI) that do not explicitly feature the guise parameter are to be read in the way just stated.

And now we see how the justification that was offered in support of (INI) in §6.3 above justifies it successfully. Where we have assumed that the inquirer knows *Q*, in order to conclude under that assumption that it would be a waste of his time and resources to inquire into *Q* (a course of action that is counterproductive), the more specific kind of scenario that we were thinking of is one where the inquirer *explicitly* knows *Q* as presented under a certain guise *g.* That is, we have assumed that the

[78] Archer (2018) objects to Friedman's (2017) ignorance norm for interrogative attitudes on the basis of cases of the following sort: the inquirer knows *Q* at *t* even though she has forgotten *Q* at *t*, and thus inquires into *Q* at *t*, if only to jog her memory and make her knowledge active again. Since our (INI) is actually concerned with *explicit* knowledge, however, such cases would constitute no objection to it (presumably, when one knows *Q* without remembering *Q* one can only be said to know *Q* implicitly).

inquirer knows a maximal fine-grained answer *a* to *Q* as presented under *g*—so that he *already* represents the world he is in as being such that *a*, and he represents the world as being that way through some declarative rendering of *g*. Inquiring into *Q as presented under guise g* would then be a waste of his time and resources, even though it need not be a waste of his time and resources to inquire into *Q as presented under some other guise h.* For the fact that a previous gap in his picture of reality has already been filled with a declarative rendering of *g* doesn't entail that it has also been filled with a declarative rendering of *h*. Furthermore, it need not be a waste of the inquirer's time and resources to inquire into *Q* under some guise or other when he knows *Q* only implicitly. There is certainty instrumental value in making implicit things explicit.

We decide to integrate the Ignorance Norm of Inquiry (INI) into our system of norms of inquiry, then, alongside the norms (GG) and (GF). Other important objections to (INI) and the way we have justified it will be addressed in the next chapter (see especially §§7.4–7.5).

We now proceed to look at yet another candidate norm of inquiry, the Knowledge Norm of Inquiry:[79]

(KNI) One should inquire into *Q* only if one knows that *Q has a true maximal answer*.

That *Q* has a true maximal answer means that there is a *p* such that *p* is true and *p* is a maximal answer to *Q*.

As with our other norms from above, the semantic contribution of 'should' takes wide-scope in (KNI). The maximal answers to a question, remember, are the complete answers to it such that they are not proper subsets of any other complete answers to that question (in the case of polar question, those are just the *yes* and the *no* answers).[80] Where '*mQ*' denotes the set of the maximal answers to *Q*, let us then use '$K(\exists p\colon p \wedge p \in mQ)$' to translate 'one knows that *Q has a true maximal answer*'. And let '*IQ*' again translate 'one inquires into *Q*'. So the logical form of (KNI) is the following: $S(IQ \supset K(\exists p\colon p \wedge p \in mQ))$.

---

[79] The norm is named this way here after Willard-Kyle (forthcoming), who defends an analogous norm for *interrogative attitudes*. We are conceiving of our norms again as norms for *a certain kind of activity* (see §1.2), in that they concern what we should or shouldn't *do*. Our norms are not norms for or agains holding intensional attitudes, related to norms of the latter sort as they may be. Maybe in a more *indirect* manner, our norms also bear on the norms that Friedman (2017) and Willard-Kyle (forthcoming) have in mind—see §10.1 for more on this. Both Friedman and Willard-Kyle conceive of their norms as *epistemic* norms, though they also call them 'norms of inquiry'. The norms investigated here are instrumental norms or norms of practical rationality. Whether they should also be called 'epistemic' is an issue that won't occupy us here. See Thorstad (2021) for relevant discussion on this.

[80] Quite generally, p is a maximal answer to Q iff $p \in Q$ and there is no $q \in Q$ such that $p \subset q$.

To see how this norm finds fault with inquirers who are *too* ignorant, consider the logically equivalent formulation $S\neg(IQ \wedge \neg K(\exists p\colon p \wedge p \in mQ))$. That is, one shouldn't inquire into a question while ignoring that *it has a true maximal answer*. Consider for example the question of *whether God was happy or rather sad when he created the universe*. That question has two maximal answers, namely, that *God was happy when he created the universe*, and that *God was sad when he created the universe*. According to (KNI), an inquirer is doing something wrong who inquires into that question but fails to know that at least one of those two propositions is the case (say, someone who doesn't even know that *God created the universe* in the first place, which is one of that question's presuppositions—more on this in §7.2).

Does the Knowledge Norm of Inquiry also need further trimming? Should we add it to our system? We continue our investigation by pursuing these questions next.

Chapter 7

§7.1 That the Knowledge Norm of Inquiry also leaves important details implicit

The Knowledge Norm of Inquiry (KNI) tells us that one shouldn't inquire into a question unless one knows that *there is a true maximal answer to that question*. But considerations about guises or representational vehicles that carry the same coarse-grained information call for a more explicit formulation of this norm, too.

Let us use the example from §6.4 again. The question of *whether 37 is a prime number* is the same as the question of *whether π is a rational number* (the same coarse-grained question). But that same question can be presented or entertained by an inquirer under different guises or vehicles, say, 'whether 37 is a prime number' and 'whether *π* is a rational number'.[81]

Now consider an inquirer who inquires into the question of *whether π is a rational number* as presented under the guise 'whether *π* is a rational number', but who fails to know that *that question has a true maximal answer* when it is presented under that very same guise, say, because she is in doubt about whether '*π*' stands for something in the first place (that expression could lack reference as far as she can tell). So far, it looks as if our inquirer is violating (KNI), and that we can find exactly the same kind of fault with her activity of inquiry that (KNI) was designed to flag. Perhaps before pursuing her inquiry into that question under that guise, the inquirer should first make sure that '*π*' has a reference. Only then will she be normatively on the clear to inquire into *whether π is a rational number* under the guise 'whether *π* is a rational number'. When she learns that *her question has a true maximal answer*, as presented under that guise, she thereby rules out the possibility that her own activity of inquiry will be an empty pursuit (a search for a true answer that doesn't exist).

But the fact that our inquirer can entertain the very same question under a different guise seems to prevent us from concluding right away that she violates (KNI) in that type of scenario. For it might be that she knows that the question of *whether π is a rational number*—which is the same as the question of *whether 37 is a prime number*—has a true maximal answer when it is presented under the guise 'whether 37 is a prime number' (she doesn't have the same kind of doubt about '37' as she has about '*π*'). The formulation of (KNI) was, remember:

---

[81] In a nutshell, again, *v*('whether 37 is a prime number') = *v*('whether *π* is a rational number'), or *whether 37 is a prime number* = *whether π is a rational number*, though 'whether 37 is a prime number' ≠ 'whether *π* is a rational number'.

(KNI) One should inquire into *Q* only if one knows that *Q has a true maximal answer*.

And now we have found a case where someone is inquiring into a question *Q* (*whether π is a rational number*) such that she fails to know that *Q* has a true maximal answer when it is presented under one guise ('whether *π* is a rational number'), and yet she knows that *Q* has a true maximal answer when it is presented under a different guise ('whether 37 is a prime number').[82] So does she or does she not abide by (KNI)? It might be suggested that she does *and* does not abide by it—but that's a contradiction, and contradictions are false.

To circumvent the issue, then, we have to pair guises in the antecedent and the consequent of the conditional embedded in (KNI), just like we did with respect to (INI):

(KNI) One should inquire into *Q*, as presented under guise *g*, only if one knows that *Q has a true maximal answer*, as presented under *g*.

So the inquirer from above violates (KNI) when it comes to the question of *whether π is a rational number* as presented under the guise 'whether *π* is a rational number' (because she is inquiring into that question under that guise and she does *not* know that that question, as presented under that guise, has a true maximal answer), even though she abides by (KNI) when it comes to the very same question as presented under the guise 'whether 37 is a prime number' (she knows that the question as presented under the latter guise has a true maximal answer). Contradiction avoided.

What is it to know that the question, as presented under some guise or other, has a true maximal answer? It depends on the case, in the sense of what information one will thereby be in possession of. In the case of polar questions, it consists of knowing an instance of excluded middle under a guise that embeds some interrogative whose content is that question. Where *Q* is a polar question, and *p* and ¬*p* are its two maximal answers, the coarse-grained information that *Q has a maximal true answer* is the same as the coarse-grained information that *either p or ¬p*.

For example, to know that the question of *whether free will is compatible with determinism,* as presented under the guise 'Is free will is compatible with determinism?', has a true maximal answer, is to know that *either free will is compatible with determinism or it isn't* under a guise such as 'The question of whether free will is

---

[82] Of course, *Q* as presented under *g* has a true maximal answer if and only if *Q* as presented under *h* has a true maximal answer—but we have used '*Q*' within the context of a knowledge ascription in (KNI): a linguistic context that famously resists the substitution of necessarily equivalent sentences under its 'that'-clause.

compatible with determinism has a true answer' (or any other guise that embeds 'whether free will is compatible with determinism' and carries the coarse-grained information that *either free will is compatible with determinism or it isn't*).

But not all cases of knowing that a question has a true maximal answer under a certain guise, and therefore not all cases where the inquirer abides by (KNI), are cases of knowing an instance of excluded middle under that guise (which might be thought of as a relatively easy accomplishment). Consider for example the question of *who killed Rajiv Ghandi*. To know that that question, as presented under the guise 'Who killed Rajiv Ghandi?', has a true maximal answer is *not* to know an instance of excluded middle under a guise that embeds 'who killed Rajiv Ghandi'. It is rather to know that *someone killed Rajiv Ghandi* under a guise that embeds 'who killed Rajiv Ghandi' (where the coarse-grained information that *someone killed Rajiv Ghandi* is the same as the coarse-grained information that *the question of who killed Rajiv Ghandi has a true maximal answer*).

We can still expand upon that way of understanding what it is for one to know that a question has a true maximal answer when the question is presented under a particular guise. In particular, where *Q* is presented under some guise *g,* we might also count the inquirer as knowing that *Q* has a true answer as presented under *g* when the inquirer knows that fact under a guise that embeds either *g,* or some variation of *g*, or some vehicle that is similar to *g* under some relevant aspects (here we could use the notion of a similarity class *f*(*g*) such as the one from §4.2). Furthermore, the guise under which one is to have the relevant kind of knowledge should embed not only an interrogative complement of the relevant type, but it should also combine that complement with a truth-predicate or some similar device.

But we leave the exploration of such details for future extensions of this work.

§7.2 That the Knowledge Norm of Inquiry entails the No False Presupposition Norm and the Anti-Dissonance Norm

In the previous section, we fleshed out a more carefully formulated version of the Knowledge Norm of Inquiry—one that avoids problems that would pester its more unqualified version (guise problems). We now proceed to explore some of the consequences of that norm. In particular, we look at some norms of inquiry that follow from it. In doing that, we omit the now expected pairing of guises from the formulation of our norms, in the interest of avoiding clutter (the pairing of guises is supposed to be there, though invisible).

We start by noting that (KNI) entails the following norm, which we will call the 'Anti-Dissonance Norm':[83]

(ADN) One shouldn't inquire into *Q* while knowing that *Q does not have a true maximal answer*.

We call the norm that way because an inquirer who keeps looking for the true answer to a question despite the fact that she knows that *there isn't a true maximal answer to it* seems to be at odds with herself (at least when the question is presented under the same guise in both cases). Since she knows that *the question doesn't have a true maximal answer*, why search for one? That sounds cognitively dissonant.

(KNI) entails (ADN). The fact that one knows that ¬*p* entails that one doesn't know that *p*. So the fact that one knows that *Q does not have a true maximal answer* entails that one doesn't know that *Q has a true maximal answer.* If (KNI) is true, then, one shouldn't inquire into *Q* while knowing that *Q does not have a true maximal answer* (ADN). For, if one were to inquire into *Q* while knowing that *Q does not have a true maximal answer*, one would be inquiring into *Q* while not knowing that *Q has a true maximal answer*—and that is exactly what (KNI) says one shouldn't do.[84]

But (ADN) does not entail (KNI). In general, the fact that one doesn't know that *p* does not entail that one knows that ¬*p*. In particular, the fact that one doesn't know that *Q has a true maximal answer* does not entail that one knows that *Q does not have a true maximal answer*. One might, for example, suspend judgment about *whether Q has a true maximal answer*, in which case one neither knows that *Q has a true maximal answer* nor knows that *Q does not have a true maximal answer*.

So (ADN) is strictly *weaker* than (KNI). Every activity of inquiry that is reproached by (ADN) is also reproached by (KNI), but not every activity of inquiry that is reproached by (KNI) is also reproached by (ADN).

Consider, for example, someone who is inquiring into *when Socrates wrote the Critique of Pure Reason* while knowing, as we do, that *Socrates did not write the Critique of Pure Reason*, and therefore knows that *that question does not have a true maximal answer*. It is not only (ADN) that issues a reproach in this case, but also (KNI). Our inquirer is not only a *dissonant* inquirer (for inquiring into a question she knows

---

[83] Relatedly, see also Friedman (2017, pp. 315–316) who points out that, in a situation where the inquirer realizes that *Q* has a false presupposition, 'further inquiry into *Q* would be irrational or epistemically inappropriate'. See Willard-Kyle (forthcoming) for a norm like (ADN) for interrogative attitudes.

[84] Quite generally, where *φ* entails *ψ*, *S*¬*ψ* entails *S*¬*φ*. Now just substitute 'one inquires into *Q* and knows that *Q* does not have a true answer' for '*φ*' and 'one inquires into *Q* and doesn't know that *Q* has a true maximal answer' for '*ψ*'.

there is no true answer to), but also one that *ignores too much* (she fails to know that the question she is inquiring into has a true maximal answer).

In contrast, consider a case where someone is inquiring into *where in Greece Socrates was born*, but who suspends judgment about *whether Socrates was born in Greece* (she can't rule out, say, the possibility that Socrates was born in the region that is now part of Syria). Since our inquirer doesn't know that *the target question has a true maximal answer*, (KNI) finds fault with her activity of inquiry. But since she also doesn't know that *the target question does not have a true maximal answer* (it might have one as far as she can tell), (ADN) does *not* find fault her activity of inquiry.[85] The Knowledge Norm of Inquiry (KNI) makes the life of an inquirer harder than just the Anti-Dissonance Norm (ADN).

(KNI) also entails the following norm, the No False Presupposition Norm:

(NFP) One shouldn't inquire into *Q* if *Q* has a false presupposition.

For, if a question has a false presupposition, then it doesn't have a true answer—so the inquirer cannot know that *it has a true maximal answer* (knowledge is factive). Since (KNI) reproaches inquiry into any question such that the inquirer doesn't know that *it has a true maximal answers*, it reproaches inquiry into questions that rely on false presuppositions. That is, (KNI) entails (NFP).[86]

The notion of presupposition that we are using in this case is again a semantic one (already introduced in §2.4—to be distinguished from the notion of the contextual presuppositions shared by speakers discussed in §§5.3–5.4). A question *Q* presupposes that *p* in this sense just in case *p* is a necessary condition for the truth of any of the complete answers to *Q*.[87] That is, a question *Q* presupposes that *p* if and only if, for any $a \in Q$, necessarily, *a* is the case only if *p* is the case.

For example, the question of *whether you're still polishing your dog's nails* presupposes that *you have a dog*. The complete answers to that question are those that entail that *yes, you keep polishing your dog's nails*, on the one hand, and those that entail that *no, you don't keep polishing your dog's nails* (you stopped doing that). No matter which one of these answers we pick, they entail that *you polished your dog's*

---

[85] Willard-Kyle (forthcoming) takes such cases to show that a norm like (ADN) for the attitude of wondering is not enough to establish what epistemically proper inquiry is, and that (KNI) is needed for that purpose.

[86] The same pattern as the one from footnote 4 applies here again. Only now we substitute 'one inquires into *Q* and *Q* has a false presupposition' for '*φ*' and 'one inquires into *Q* and doesn't know that *Q has a true maximal answer*' for '*ψ*'.

[87] See also Belnap and Steel (1976, p. 5), van Fraassen (1980, p. 140), Groenendijk and Stokhof (1984, pp. 31-32).

nails before. There is no way for that question to have a true answer without you having polished your dog's nails before, *ergo* that question presupposes that *you polished your dog's nails before*. And, hopefully, that is false, so that the target question relies on a false presupposition.

We just saw that the Knowledge Norm of Inquiry (KNI) entails the No False Presupposition Norm (NFP). But not the other way around—(NFP) does not entail (KNI). For the fact that *one doesn't know that Q has a true maximal answer* does not entail that *Q has a false presupposition*. One might fail to know that *Q has a true maximal answer* for other reasons than that *Q* has false presuppositions, and therefore no true answers.

So (NFP) is logically related to (KNI) in the same way that the Anti-Dissonance Norm (ADN) is related to (KNI). At this point it looks as if (NFP) and (ADN) are on the 'same level' in the logical order of norms of inquiry, seeing as both are entailed by (KNI), though neither of them entails (KNI).

But actually (NFP) entails (ADN). For if one knows that *Q does not have a true answer*, then *Q* does not have a true answer—and therefore *Q* relies on a false presupposition (at least if we are reasoning classically). Since (NFP) finds fault with inquiry into a question that relies on a false presupposition, and since inquiry into a question one knows there isn't an answer to is inquiry into a question that relies on a false presupposition, (NFP) also finds fault with inquiry into a question one knows there isn't an answer to.

So the right logical order of these norms is: (KNI) $\models$ (NFP) $\models$ (ADN), where '$\models$' expresses the entailment relation. In other words, the Knowledge Norm of Inquiry (KNI) is strictly stronger than the No False Presupposition Norm (NFP), which is in turn strictly stronger than the Anti-Dissonance Norm (ADN).

§7.3 Where we become doubtful about including the Knowledge Norm of Inquiry into our system of norms

At least in a large class of cases, the Knowledge Norm of Inquiry (KNI) issues reproaches where reproaches are called for. Consider, for example, someone who is inquiring into the *how one can prove the fact that the Earth is flat*. Since that inquirer doesn't know that the target question has a true answer, (KNI) rightly says that she is doing something wrong.

But so does the weaker norm of No False Presupposition (NFP), seeing as that question presupposes that *the Earth is flat*, which is false (we are assuming that *one can prove the fact that p* only if *p*). Maybe, then, we don't need (KNI) in our system of norms—we can just include (NFP) in it, and thereby also get the Anti-Dissonance Norm

(ADN) as a consequence of it. That would be the right theoretical choice for us to make, were we to find cases where (KNI) issues a reproach though a reproach is *not* called for. In order to put that to a test, remember, we are thinking of the 'should' of our norms as the 'should' of instrumental or practical rationality. If lack of knowledge that *there is a true maximal answer to Q* doesn't render inquiry into *Q* into a counterproductive activity vis-a-vis the very goal of the activity of inquiring into *Q*, and neither is knowledge of the presence of such an answer a means to the goal of the activity of inquiring into *Q*, then (KNI) is false—for (KNI) is advanced here as a universal instrumental norm.

Suppose a physicist—call him 'Physicist'—is inquiring into *how the universe began*, maybe under a guise such as 'Did the universe begin with the Big Bang or in some other way?' (but let us not worry about guises from now on, and just concentrate on pure content or the coarse-grained questions themselves). Physicist does what physicists do when inquiring into such questions, such as trying to find the best explanation for facts concerning the relationship between the distance and speed of astronomical objects (apparent expansion of the universe), etc. All the while, however, even as Physicist is inquiring into that question, a more fundamental question remains unsettled for him, namely, the more fundamental question of *whether the universe had a beginning* to begin with.

Any complete answer to the question of *how the universe began* entails that *the universe had a beginning*—the latter is a presupposition of the former. But Physicist suspends judgment about *whether the universe had a beginning*, and suspends judgment accordingly about *whether the question of how the universe began has a true answer*. For all he knows, there might not be one. So he doesn't know that *the question of how the universe began has a true answer*, and therefore he doesn't know that *the question of how the universe began has a true maximal answer*.

But it doesn't seem that Physicist must be at fault in inquiring into the question of *how the universe began* in such a case, despite the fact that he does not know that *that question has a true maximal answer*. At leas not from a purely instrumental point of view. Physicist's inquiry into the target question isn't made counterproductive—or no instrumental value is detracted from it—just by virtue of the fact that he fails to know that the target question has a true answer.

In fact, if the question of *how the universe began* does have a true answer, unbeknownst to Physicist as that fact may be, then not only may he eventually come to know that answer—he may thereby also settle his other, more fundamental question of *whether the universe had a beginning*. For any complete answer to the former question entails the *yes* answer to the latter question. His settling the question of *how the universe began* is a possible means to his settling his other question of *whether the*

universe had a beginning. As far as our assumptions go, our physicist's activity of inquiry into the former question might be rather useful or conducive to its very goal, as well as the goal of settling other related questions he is still undecided about. The very fact that *there is* a true complete answer to his question means that his inquiry is on the right track, in fit with reality (he is on to the facts), whether he knows this or not.

On the contrary assumption that all answers to the question *of how the universe began* are false, in which case *there isn't* a true answer to that question, instrumental value will indeed be detracted from his activity. For then he will *never* come to know what that question's true answer is—simply because *there isn't* one. But then the No False Presupposition Norm (NFP) does the work of telling us what he is doing wrong (if all the complete answers to the question are false, then the question relies on a false presupposition), and we have no need to appeal to the Knowledge Norm of Inquiry (KNI) to explain the error of his ways.

It is rather the *absence of a true answer* that makes inquiry into the question counterproductive here—not the *lack of knowledge that there is one*. These considerations suggest that the inquirer himself doesn't have to rule out the possibility that his own activity of inquiry will be an empty pursuit (a search for true answers where there is none to be had). It just shouldn't *be* an empty pursuit.

But matters are not that straightforward. The following example from Willard-Kyle (forthcoming) seems to support the idea that (NFP) isn't yet strong enough to properly assess certain activities of inquiry, and that we need (KNI) to precisely pinpoint where they go wrong. The inquirer knows that *one of Mercury, Venus, Earth or Mars is the largest planet of our solar system*, though she doesn't know *which one of those four planets is the largest one*. She then goes on to inquire into the question of *whether it is Earth or rather Mars that is the largest planet of our solar system*. She doesn't know that *that question has a true answer*, for she doesn't know that *it is either Earth or Mars that is the largest planet of our solar system*. She is leaving Mercury and Venus out as candidates for the largest planet. It seems that her activity of inquiry is going wrong.

Nevertheless, given that the Earth is indeed the largest planet of our solar system, she is not violating the No False Presupposition Norm (NFP). The question of *whether it is Earth or rather Mars that is the largest planet of our solar system* does not rely on a false presupposition (it is after all true that *either the Earth is the largest planet or Mars is the largest planet in our solar system*). But she *is* violating the Knowledge Norm of Inquiry (KNI)—and that is what she is doing wrong.

There are ways to resist this argument in support of (KNI) understood as an instrumental norm, however. The reply would be that such an activity of inquiry is *not* going wrong, all the appearances to the contrary notwithstanding, at least not *qua* an activity devoted to the goal of finding the true answer to the question *whether it is Earth*

or rather Mars that is the largest planet in our solar system. For *there is* such an answer indeed, and that activity is not doomed to failure just because its actor doesn't know that. She doesn't have to be aware that she is on the right track in order for her to *be* on the right track. At least she is not instrumentally at fault in performing that activity. Her mistakes are elsewhere, and these will explain away the impression that her activity of inquiry is at fault. What could those other mistakes be?

Maybe we read off of the description of the case that the inquirer is wondering about the target question while taking it that it may not have a true maximal answer. So she would be wondering *whether it is Earth or rather Mars that is the largest planet of our solar system* while believing that *Mercury might as well be the largest planet of our solar system*, say. And that sounds like an incoherent combination of attitudes. Her problem is not that her inquisitive performance is counterproductive—for she really is on to the facts by pursuing her question, and she may still eventually reach the happy end of her inquiry (knowing the true answer to her question, namely, that *the Earth is the largest planet in our solar system*). Her problem here is rather that her intensional attitudes fail to cohere with each other.

Notice, furthermore, that even if we suppose that our inquirer believes that the question she is wondering about here *does* have a true maximal answer, we are still going to find grounds to criticize her attitudes—though this time on account of the fact that her belief is not supported by what she knows or by what her evidence is. For, again, as far as what she knows goes, not only the Earth and Mars, but also Mercury and Venus could be the largest planet of the solar system. So she is not epistemically entitled to believe that *it is either the Earth of Mars that is the largest planet of our solar system*.

Either our inquirer's attitudes harbor incoherence, or they fail to be supported by her knowledge or evidence, then.[88] In neither case, however, is her goal-oriented performance of looking for the true answer to her question turned counterproductive by virtue of the fact that she doesn't know that *her question has a true answer*.

Lack of knowledge that a question has a true maximal answer is not detrimental to the success of inquiring into that question, though the sheer lack of a true maximal answer is. Perhaps knowledge that a question has a true maximal answer can still be instrumentally valuable to the activity of inquiring into that question, say, by helping the inquirer implement the proper means to the goal of knowing the question's true maximal answer (when something is instrumentally valuable to an activity because it is

---

[88] These would be failures of *structural* rationality and *substantial* rationality, respectively. For the idea that all failures of the former kind boil down to failures of the latter kind, see Kolodny (2005, 2007). See Worsnip (2021) for the view that these are independent dimensions of rationality. See Rosa (2022) for an alternative to both of these views.

a means to the means to the goal of that activity). Perhaps such knowledge will make the inquirer more motivated and resilient to challenges in pursuing her inquiry, compared to an inquirer who doesn't yet rule out the possibility that there isn't a true maximal answer to her question.

Perhaps—but here the connection between the kind of knowledge that (KNI) requires for proper inquiry and the maximization of the goals of inquiry becomes far too indirect and flimsy, and therefore not strong enough to vindicate (KNI) as a universal instrumental norm of inquiry.

§7.4 That inquiry into *Q* is aimed at knowing *Q*

So we have found good grounds for excluding the Knowledge Norm of Inquiry (KNI) from our system of norms. Luckily, however, in the next chapter (§8.2) we will formulate a norm of inquiry that *also* entails the No False Presuppositions Norm (NFP), though it doesn't entail the overly strong (KNI). The new norm tells the inquirer that it should at least be *possible* for her to know the true maximal answer to the question she is inquiring into.

Before moving onto that, however, it is time for us to step back and look more carefully at the way in which we have been justifying our acceptance and rejection of these norms of inquiry.

In particular, in justifying our acceptance or rejection of purported norms of inquiry, we have relied on the assumption that the activity of inquiry into *Q* (inquiring into *whether it rains*, inquiry into *who did it*, etc.) constitutively involves the aim or goal of knowing *Q* (knowing *whether it rains*, knowing *who did it*, etc.). Inquirers want or need to know the answers to the questions they are inquiring into—that is the kind of want or need that the activity of inquiry is supposed to satisfy.

To say that knowing *Q* is a *constitutive* goal of the activity of inquiring into *Q* is to say that a cognitive agent who doesn't want or need to know *Q* is not really inquiring into *Q*, possible appearances to the contrary notwithstanding (perhaps he is faking, or inquiring into some other related question that is not identical to *Q*, etc.). As should be expected by now, a pairing of guises is to be read off of this claim about inquiry too. It is not only that the inquirer wants or needs to know *Q*—she wants or needs to know *Q under a certain guise* (presumably because knowledge of *Q* under the relevant guise is expected to play a certain role in her mental economy).

We have been relying heavily on this view about the nature of inquiry. For example, in justifying the Ignorance Norm of Inquiry (INI) above we have said: given that the inquirer already knows *Q* (under guise *g*), both information-gathering and armchair inquiry into *Q* (as presented under *g*) will be counterproductive activities for her to engage with. For, in that case, the constitutive goal of those activities has

already been achieved. Similarly, we offer the following justification in support of the No False Presupposition Norm (NFP): assuming that *Q* has a false presupposition in the world *w* the inquirer is in, it follows that every complete answer to *Q* is false in *w*, and therefore inquiry into *Q* is a counterproductive activity for that inquirer to perform in *w*. For that inquirer won't *ever* achieve the constitutive goal of that activity in that world—she won't ever know *Q* at *w*.

The idea that knowledge is the aim of inquiry is widely endorsed in the literature.[89] But of course talk of 'the goal of inquiry' can be misleading, on more than one count. First, it is *people* and agents more generally who literally speaking have goals. Inquiry is a type of activity. Both this abstract type and its tokens can only be said to 'have goals' as they inherit the goals of the actors who execute or instantiate them. Neither the type of activity nor its tokens can be *literally* said to have goals, or to aim at something.

Second, the actors of activities of inquiry might have many different goals in performing them. They might engage in it because it is fun, or because they want to be paid at the end of it, for example. The claim that inquiry into a question is aimed at knowing its true answer is presumably not supposed to exclude such possibilities.

Note, however, that there is a contrast between these other goals that one might have when engaging in inquiry (fun, money, etc.) and the goal of knowing the answer to the question. To bring out the contrast, let us consider speech acts whereby a speaker makes it manifest that she is inquiring into a question, and she does that by sincerely uttering an interrogative sentences. For example, a speaker may express her inquisitive thoughts about *whether virtue can be taught* by sincerely uttering 'Can virtue be taught?'. Now compare utterances such as:

(1) Does my daughter hate me? It is not fun to think about this.
(2) Do you know what time Constanza left the office? I'd appreciate your help, for I am not even getting paid to look into this.

with:

(3) # When are you going to call me? I neither want nor need to know when you are going to call me.
(4) # Are quarks divisible? I don't care if I know the answer.

[89] See for example Schaffer (2005), Whitcomb (2010), Friedman (2017).

Surely there is a bit of dissonance in *all* of (1)–(4). The speaker who utters (1), for example, is trying to figure something out despite the fact that this is admittedly unpleasant to him/her. And we sense a complaining mood about having to inquire in the utterance of (2), too. (3) and (4) also signal dissonance[90]—but there is a significant difference between (3)/(4), on the one hand, and (1)/(2), on the other. The difference is that it is hard to imagine the utterers of (3)/(4) *sincerely* uttering their interrogatives, assuming that what they said in the second bit of their utterance is true.

The thesis that inquirers always aim at knowing the answers to their questions explains the relevant contrast: the speaker cannot be speaking truly while inquiring *via* asking in (3) and (4). (Asking is a way of inquiring). Wanting or needing to know the answer to a question is constitutive of inquiry into that question. So if the speaker is really inquiring into the question expressed by the interrogative in (3), for example, then what she says in the second bit of her utterance is false: she *does* want or need to know whether her addressee is going to call her. Contra-positively, if what she says in the second bit of her utterance is true, then it seems that her utterance of the interrogative 'When are you going to call me?' is not an act of inquiry after all. Why would she be inquiring into that, if she doesn't want or need to know the answer to the target question? Similar points apply to the utterance of (4).

In contrast, surely an inquirer can perform inquiries without aiming for fun or money. The desire or need for fun/money is not constitutive of inquiry. That explains why (1)/(2) are not as problematic or odd-sounding as (3)/(4) are. The utterer of (1)/(2) can perfectly well be inquiring through the first bit of her utterance (the interrogative bit) *and* saying the truth in the second bit (the declarative bit).

There is the explanation, and then there is the explanation of the explanation. In our case, our explanation of the oddity of utterances such as (3)/(4) is that knowing *Q* is a constitutive goal of the activity of inquiring into *Q*. The explanation of that explanation would be an explanation why someone who inquires into *Q* thereby aims at knowing *Q*. As already suggested above, presumably the inquirer wants or needs to know *Q* (under a certain guise) because that kind of knowledge is expected to play a certain role in her cognitive economy, such as that of allowing her to make decisions and draw inferences, or furthering some of her other goals, etc. Inquiry is her way of trying to satisfy such a need or want—that is what the kind of creature that she is has to do in order to achieve that goal in the world she is in.

An empirical hypothesis here, then, is that the practice of inquiry came into being in order to allow us to acquire knowledge when we need it or want it (think now of

---

[90] Similarly, van Elswyk and Sapir (2021, p. 5854) note the defectiveness of:
(2) Even though I wonder who was at the party, I don't want to know who was there.

inquiry as a behavioral/cognitive trait of natural cognitive systems such as ourselves). Knowledge is a constitutive goal of inquiry because inquiry was naturally designed to be a means to knowledge when we have a need or desire for it.

We won't try to go beyond such speculative remarks here, for our methods are not the most suited to address the question (the question, that is, of *why* is it the case that inquirers want to know or have a need to know). But phenomena such as the contrast between (1)/(2), on the one hand, and (3)/(4), on the other, suggest that there is a fact to be explained here indeed—a fact whereby sense is made of the target contrast. That fact, we suggested, is the fact that knowing *Q* is a constitutive goal of the activity of inquiring into *Q*. We can stick to that explanation without yet having an explanation of that explanation.

§7.5 That inquiry into *Q* is aimed at knowing *Q* (continued)

There are yet other reasons to think that knowledge is the constitutive aim of inquiry.[91]

For example, most likely someone in the year 2023, being aware that he won't ever know the answer to the question of *what the most fashionable idea in philosophy will be in the year 3402*, will not earnestly inquire into that question either. And constitutive goals of activities have exactly this feature: typically, when an agent becomes aware that the goal is *unachievable*, he will avoid engaging in the activity.

Or consider this. Most likely someone who is aware that he already knows the answer to the question of *whether every German is blonde* will not earnestly inquire into that question either (that person already knows some Germans who are not blonde). And constitutive goals of activities also have this feature: typically, when an agent becomes aware that he has *already achieved* the goal, he will avoid engaging in the activity—for the activity becomes then pointless from his own perspective.

Good reasons to endorse the view that knowledge is the constitutive aim of inquiry are not wanting. What may seem wanting, however, is a good response to the following concern about this view: what could be wrong about inquiring into a question one already knows the answer to (constitutive goal already achieved), when all one is trying to do is to *make one's knowledge more solid*?

By an inquirer's 'making her knowledge more solid' we can mean at least two things. We can mean the inquirer's *double-checking* whether her previous answer to her question is the right one and reinforcing her belief in it as a result, or her acquiring *more or better reasons* in support of her knowledge of the question's true answer.

---

[91] See for example the case that Friedman mounts for it in her (2019, pp. 300-302). See also Kelp (2018) for persuasive arguments for the conclusion that knowledge, as opposed to justified belief or true belief or justified true belief, is the goal of inquiry.

Suppose that the inquirer already knows *whether she got the job* by knowing that *she got the job*. But she wants to make it really sure that *she got the job*, and so she starts going through the reasons that she initially had to think that *she got the job*, thus trying to determine yet again whether that is the case (on the basis of those very reasons). She is double-checking *whether she got the job*.

The worry concerning the thesis that knowledge is the constitutive aim of inquiry here is that, in cases such as this one, it doesn't look like a reproach of the inquirer's performance is called for. It is not as if our inquirer shouldn't be inquiring into *whether she got the job*, just because she already knows that *she got the job*. She knows it, but she still wants to make it sure—so it seems just right for her to inquire into the matter again. But reproach *would* be called for if knowledge were the aim of inquiry (using the norms of inquiry that we have been grounding on that thesis).[92]

We can change the case slightly in such a way that, instead of trying to determine yet again *whether she got the job* on the basis of information she already had, what the inquirer does is gather new information that bears on the issue. She is then looking for *additional reasons* (in addition to the ones she already has) to believe that *she got the job*. In either of these cases, we might say, the inquirer already knows *Q*, but she goes on to inquire into *Q* in order to make her knowledge of *Q* even more solid (armchair inquiry in the former case, information-gathering inquiry in the latter). These are supposed to be legitimate cases of inquiry without an aim to know.[93]

Their apparent legitimacy notwithstanding, however, we can challenge the conclusion that such cases feature an inquirer who inquires into *Q* without aiming to know *Q* (if they are to be *possible* cases at all).

First, we might point out that, if the inquirer is really inquiring into *Q*, then she is treating the question *Q* as an *open* question, as opposed to treating it as closed. Therefore she doesn't know *Q* at the time she is inquiring into it—she doesn't even so much as *believe* any of *Q's* complete answers during the time she is inquiring into *Q*.[94] So we can perfectly construe of the double-checker's situation, for example, as one where her belief in *Q*'s answer gets suspended during her double-checking (which *is*

---

[92] See Falbo (2023) and Woodard (forthcoming).

[93] Falbo (2023) argues on the basis of such examples that knowledge is not the aim of inquiry. She also presents a case where the inquirer inquires into *whether it will rain later* while knowing that *he is not in a position to know whether it will rain later*. From that, Falbo concludes that the inquirer is not aiming to know *whether it will rain later,* even though he is inquiring into that very question. But the latter doesn't follow from the construction of the case: the inquirer/his activity might still be aimed at knowing *whether it will rain* even though he knows that *he is not in a position to know whether it will rain*. Furthermore, the target inquirer might rather be inquiring into/wanting to know *whether it is more likely that it will rain later than not*. We will get back to this in §8.5.

[94] See Friedman (2017, 2019) for this line of thought.

an activity of inquiry), and therefore she doesn't know *Q* *while* she is engaged in that activity (even if she is to restore that knowledge after she is done with the double-checking). Her inquiry does aim at knowledge after all, and she is not inquiring into *Q* while knowing *Q*.

Although that sounds like a good response when it comes to the double-checker's case, it is less convincing when it comes to the case of the inquirer who is simply looking for additional reasons in support of her knowledge of *Q*'s answer. She need not suspend belief just because she wants to confirm her belief even further.

Our second point, however, is that we can make perfect sense of the inquirer's behavior here by taking her to be inquiring into such questions as *whether she can find more reasons to back up her knowledge that p*, or the question of *whether there are other facts that support p*, or similar questions (where *p* is the maximal answer to *Q* such that the inquirer already knows that *p*). Maybe *that* is what the inquirer wants or needs to know, but doesn't know yet, and the reason why she engaging in inquiry. Her knowledge of *Q* remains intact, but she is not exactly inquiring into *Q*—she is rather inquiring into questions concerning how other facts bear on *Q*. Her speech behavior might confirm our hypothesis, say, if in response to the question of what she is trying to accomplish she says 'I'm looking for more evidence for the fact that *p*' or something along these lines.

Depending on the nature of the case, in challenging alleged cases of inquiry into *Q* that are not aimed at knowing *Q* in this way, we need not even take the inquirer to be inquiring into such higher-order issues as *which other facts support p*. They might simply be inquiring into whether some *other propositions* are true which are logically/probabilistically connected to *p*, thereby showing sensitivity to such connections, without necessarily thinking *about* them. For example, in the case from above, where we take the inquirer to already know that *she got the job*, she might go on and pose a question to Mr. Smith from human resources by using the interrogative 'The job is mine, right?'. But what she is thereby inquiring into is the question of *whether Mr. Smith will corroborate that she got the job* (she is craving a feeling of reassurance), which is different from the question of *whether she got the job*. This reply brings to our attention a fact that we all have to bear with—even independently of this dispute about the constitutive aim of inquiry—namely, that which questions an inquirer is inquiring into is often left underdetermined by her behavior.

Finally, maybe the description of the target cases is perfectly accurate: the inquirer *is* inquiring into *Q* and she knows that *Q* at the same time. But whence does it follow that she is not aiming to know *Q*? She might know *Q* without knowing that she

knows *Q*, for example.[95] She is aiming to accomplish something that she is not aware she has already accomplished. From her perspective, maybe she still doesn't know the answer—so inquire she will. In general, agents may aim at bringing things about that have already been brought about. Just the fact that the inquirer already knows *Q* (perhaps implicitly) doesn't guarantee that in performing her activity of inquiry she is not aiming at knowing *Q*.

We conclude that the types of cases presented above (cases involving double-checking and the search for additional reasons) do not show that knowledge is not a constitutive aim of inquiry. We hold on to our thesis that knowledge is a constitutive aim of inquiry, then, and we continue to build our system of norms on the basis of it.

[95] Of course, this contradicts the so-called 'KK Principle'—see Greco (2014) for a recent defense of this principle.

Chapter 8

§8.1 That our project wouldn't need to be completely abandoned were we to take the settling of a question *Q* to be more, or less, than knowing *Q*

Our initial characterization of inquiry from §1.1 was that it was an activity aimed at settling a question. We later saw in §2.1 that settling a question must be some form of epistemic success. Otherwise, we wouldn't be able to distinguish the questions that the inquirer is already in a position to settle from the questions that she is not yet in a position to settle in the way we wanted.

In the last two chapters, we have been relying on the view that inquiry constitutively involves the aim *to know*, and we have been defending that view against objections and grounding our system of norms on it, together with our framework for theorizing about inquiry (in terms of coarse-grained information, which are sets of possible worlds, and fine-grained information, which are pairs of those sets and representational vehicles).

Putting everything together, then, we are naturally led to the view that to settle a question *Q* is to know *Q*, thus strengthening the things we have said about the notion of *settling a question* in Chapter 2. That one has settled the question of *who won the match* entails (among other things) that one has come to know *who won the match*, or that one has learned *who won the match*. Now given that the state of affairs that the inquirer wants or needs to bring about is one where she has settled some question *Q*, it follows that the state of affairs that the inquirer wants or needs to bring about is one where she knows *Q*. So if we treat '… is a constitutive aim of…' as a transparent or non-opaque context, our view that knowledge is a constitutive aim of inquiry follows from the view that settling a question is a constitutive aim of inquiry.

If, however, the reader thinks that we have misstepped somewhere along the way—because they think that settling a question should be understood as something that doesn't require knowledge—they don't have to completely part ways with our project. They can instead reformulate our system of norms by substituting the verb 'settles' for the verb 'knows' whenever we use the latter in formulating those norms. They can then endorse the modified norms and justify them in the way we have justified them (using the same theoretical constructs).

Consider for example the Ignorance Norm of Inquiry (INI), that is, that one shouldn't inquire into *Q* if one already knows *Q*. Where the fact that the inquirer *has settled Q* entails not only that her coarse-grained information is a member of *Q*, but

also that she is in possession of a fine grained-answer to *Q*, then a norm that is just like (INI), except that it deploys 'has settled *Q*' instead of 'knows *Q*', will receive the same kind of justification that we have offered in support of (INI) above. For both information-gathering and armchair inquiry into *Q* would be counterproductive activities for an inquirer to perform when that inquirer *has already settled Q*.

That is not to say, of course, that *every* alternative interpretation that a reader could attach to the success term 'settle' would work that way. Consider for example the No False Presupposition Norm (NFP), which directs inquirers to avoid inquiring into questions with false presuppositions. We have justified that norm as follows: where *Q* relies on a false presupposition, it admits only of false answers, and no true answers—in which case the inquirer's search for the true answer to *Q* will be a counterproductive wild-goose chase. But, of course, if what the inquirer wants or needs is to settle *Q*, and settling *Q* does not even involve believing/accepting the *true* answer to *Q*, then this kind of justification for (NFP) will not work. For the justification wrongly assumes that the inquirer is searching for the true answer to *Q*, and therefore her activity of inquiry into *Q* is bound to fail to achieve its goal when there is no true answer to *Q*.

But apart from such weak, perhaps ultimately unfruitful interpretations of 'settle' (understood as a success term), our efforts at building models of inquiry, the distinctions that we have been drawing on the matter, and our system of norms of inquiry (our whole theory of inquiry and its theoretical constructs), won't amount to fruitless efforts even for those who hesitate to identify the inquirer's having settled a question *Q* with her knowing *Q*. For they can again use a rewriting mechanism such as the one suggested above to obtain similar results to the ones we obtain here.

§8.2 Where the Anti–Impossibility Norm is introduced to flag cases of misguided inquiry

In §7.3 we decided not to include the Knowledge Norm of Inquiry (KNI) in our system of norms of inquiry (one should inquiry into *Q* only if one knows that *Q has a true maximal answer*). That norm has the virtue of being *strong*, in that it entails again the No Presupposition Norm (NFP), and consequently the Anti-Dissonance Norm (ADN), both of which we found fit to include in our system. But the strength of the Knowledge Norm of Inquiry was our very reason for its dismissal: it issues reproaches where reproaches are not called for. Its theoretical virtue is also its theoretical vice.

But we promised then to formulate a norm of inquiry that *also* entails the No False Presupposition Norm (NFP), though it doesn't entail the Knowledge Norm of Inquiry (KNI) (§7.4). If we accomplish that much, we get to tighten up the boundaries of proper inquiry even more, locating the instrumentally ideal inquirer within a smaller

spectrum than the one that was left open by the prohibitive norms (NFP) and (INI), on the one hand, and the proactive norms (GG) and (GF), on the other.

Consider questions *Q* such that it is impossible for the inquirer to know *Q*. In the preface of the *Critique of Pure Reason*, Kant has some such questions in mind when he writes that reason is bound to pose itself questions that it cannot deliver an answer to.[96] Given that inquiry aims at knowledge, an inquirer shouldn't inquire into such questions. For she would then be trying to accomplish the impossible, and no one should try to accomplish the impossible (that too is a misuse of one's time and resources, seeing as it never brings about what it is supposed to bring about).

The norm we are thinking of here is what we will call the 'Anti-Impossibility Norm':

(AIN) One shouldn't inquire into *Q* if it is impossible for one to know *Q*.

As usual, we give wide-scope to 'should' here, too. (AIN) is equivalent, then, to the claim that one should be such that: one inquires into *Q* only if it is possible for one to know *Q*. A question *Q* such that it is impossible for one to know *Q* is a question that is *unsettable* or *unanswerable* to one. So (AIN) tells inquirers to avoid inquiring into questions that are unsettable or unanswerable to them.

The concept of *possibility* at play in (AIN) can be made to be more or less inclusive, other things being equal. The more scenarios count as possible in the relevant sense, the less things count as impossible for the inquirer to know. The less scenarios count as possible in the relevant sense, the more things count as impossible for the inquirer to know. A version of (AIN) featuring the former concept of possibility is then *weaker* than a version of it featuring the latter concept of possibility. The more things are impossible for the inquirer to know, the less questions will her be in the clear to inquire into according to (AIN).

Consider for example a *logical* notion of possibility. Say that a logically possible world is one such that the truths that hold in it are not inconsistent with the logical truths. Now say that 'It is logically possible that *p*' is true at a world *w* just in case there is a logically possible world where '*p*' is true (with the meaning that '*p*' has in *our* mouths in both cases). That is a very inclusive notion of possibility. The same things are logically possible relative to any possible world.

There is, for example, a logically possible world where you know *whether the Italian language will be spoken for a longer time than the Portuguese language*. But, in a good sense of these words, you *cannot* know what the true answer to that question is

---

[96] See Kant ([1781, 1787] 1998), Avii–Aviii.

—for you have no way of making such a powerful prediction about the future at the present time, no way of telling how long these two languages are going to last for, etc. Since it is logically possible for you to know *whether the Italian language will be spoken for a longer time than the Portuguese language*, however, (AIN) won't issue any reproach to your inquiry into that question under the present interpretation of the notion of possibility.

Contrast that to *circumstantial* notions of possibility, where what counts as possible changes from one possible world to the next.[97] Say that 'It is possible that *p*' or 'Possibly *p*' is true at a world *w* in this sense when there is a possible world where '*p*' is true among the possible worlds that are not ruled out by certain contingent facts of *w* (with the meanings that '*p*' and 'Possibly *p*' have in our mouths). We think of the context where those sentences are uttered as providing a set of true propositions that are assumed/known to be true by the interlocutors in the conversational background where those utterances are made, which then determines a modal base—some set of possible worlds—relative to which we assess necessity and possibility claims.[98]

So *which facts* more exactly are supposed to rule out possible worlds will depend on which truths we know or take for granted in the context where we make our utterance. Where the relevant facts include more than just logical truths, however, the notion of circumstantial possibility will be less inclusive (often *much less* inclusive) than the notion of logical possibility.

It is impossible in this circumstantial sense for you to know now *what I was thinking about on January 6, 1991 at 3 p.m*, for example (think of that impossibility claim as being asserted by me). There is no possible world where you know the true maximal answer to that question—at least not among the possible worlds where certain propositions that I know are true (in the situation of my assertion) hold: that I have no memoirs from that time for you to consult, that I myself don't remember what I was thinking, that neither of us is able to time-travel, etc. We might also want to put the point as follows: there is no *possible future* of the world we're in, relative to the present time, where you come to know *what I was thinking about on January 6, 1991 at 3 p.m*.

So, where the Anti-Impossibility Norm (AIN) deploys such a notion of circumstantial possibility, it says that you shouldn't inquire into that question. And that sounds correct, seeing as your inquiring into that question would be a misuse of your time and resources. That is perhaps a gap in our picture of reality that is bound to

---

[97] See Kratzer (1981) for the notion of circumstantial or 'root' modality.

[98] On the context dependence of sentences deploying modal expressions in general, see also Stalnaker (1970) and Kratzer (1977).

remain a gap, filled by no one at no period of human history. Examples like this one multiply easily—there will always be uncountably many gaps in our picture of reality.[99]

Note, however, that if you so much as inquire into the question of *what I was thinking about on January 6, 1991 at 3 p.m.*, neither the No False Presupposition Norm (NFP) nor the Ignorance Norm of Inquiry (INI) will flag the error of your ways. Since *I exist*, and I was alive and thinking about something indeed at the relevant time and date, the question doesn't rely on false presuppositions. So you don't violate (NFP) when you inquire into that question. And in doing that you don't violate (INI) either, for surely you don't already know the true maximal answer to that question.

We have a case, then, for including the Anti-Impossibility Norm (AIN) in our system of norms, even where it features a stricter notion of possibility such as the notion of circumstantial possibility. It is counterproductive for an inquirer to inquire into questions *Q* such that it is impossible for them to know *Q* in that sense—and neither of the other prohibitive norms that we have accepted so far tells the inquirer to avoid inquiring into such questions. The new norm allows us to tighten up the limits of proper inquiry even further.

§8.3 That the Anti-Impossibility Norm entails the No False Presupposition Norm, though not vice-versa

We also have a logico-explanatory reason to accept the Anti-Impossibility Norm in one of its versions, namely, that it entails some of the norms we have already accepted.

Our norm has such a property at least if its notion of possibility is a notion of circumstantial possibility—henceforth *possibility*$_c$—which satisfies the following constraint: it is possible$_c$ for an inquirer to know *Q* at *w* only if *there is* a maximal answer to *Q* that is true at *w*. So that, if some of *Q*'s presuppositions are false at world *w* (all the complete answers to *Q* are false at *w*), then it is impossible$_c$ for any inquirer to know *Q* at *w*.

That only makes sense. It is impossible for any human being in the actual world to know *why the sun revolves around the Earth*, we might say—for that question relies on the actually false presupposition that *the sun revolves around the Earth*. None of us can know *why the sun revolves around the Earth*.

To make the notion of possibility deployed in our Anti-Impossibility Norm explicit:

(AIN) One shouldn't inquire into *Q* if it is impossible$_c$ for one to know *Q*.

---

[99] For there are uncountably many questions concerning, for example, uncountably many possible magnitudes of physical objects, as measured by real numbers.

In this way, we see that (AIN) entails the No False Presupposition Norm (NFP). For (NFP) is, remember:

(NFP) One shouldn't inquire into *Q* if *Q* has a false presupposition,

which is equivalent to:

(NFP) One should inquire into *Q* only if none of *Q*'s presuppositions are false.

And (AIN) is again equivalent to:

(AIN) One should inquire into *Q* only if it is not impossible$_c$ for one to know *Q*.

And that *it is not impossible$_c$ for one to know Q* entails that *none of Q's presuppositions are false*—so that (AIN) entails (NFP). In order to abide by (AIN), the inquirer has to abide by (NFP). If she violates (NFP), she thereby violates (AIN).

We have already accepted and justified (NFP). Inquiry is, among other things, a search for true answers to questions, an attempt to settle them. So inquiry into a question with false presuppositions is a search for *what doesn't exist*. For a question with false presuppositions is a question for which there are no true answers, only false answers. Since (AIN) entails (NFP), that gives us yet another reason to accept the former (over and above the fact that it correctly classifies certain cases of inquiry as instrumentally improper).

Furthermore, we were also happy to accept the Anti-Dissonance Norm:

(ADN) One shouldn't inquire into *Q* while knowing that *Q does not have a true answer*.

We saw in §7.2 that the No False Presupposition Norm (NFP) entails the Anti-Dissonance Norm (ADN) (at least when we are reasoning classically). Now since the Anti-Impossibility Norm (AIN) entails (NFP), it also entails (ADN). The new norm (AIN) not only tightens up the boundaries of proper inquiry further—it also adds *unifying power* to our system of norms of inquiry.

Over and above cases that involve a more *local* kind of impossibility to know the answer, such as cases of questions with contingently false presuppositions, (AIN) will also flag the impropriety of inquiry into questions whose true answers are impossible for the inquirer to know in a more *global* sense.

When Kant talks about our epistemic limitations regarding the purported postulates of morality, for example—God's existence, freedom of the will, the

immortality of the soul—he is not only talking about *his* limitations in his *particular* situation regarding the questions of *whether God exists*, etc.[100] He rather has in mind a more structural kind of limitation, a limitation of *human reason*. It applies to all of us, regardless of our individual variations in cognitive skills and knowledge. It would be impossible at *any* time for *any* of us to know *whether God exists*, etc. Or consider whichever other questions none of us is in a position to settle, on account of more deeply structured limitations—e.g. essentially undecidable problems in computation theory and logic/mathematics.[101]

That (AIN) tells against searching for the true answers to such questions (even if *those* examples are not quite right) is secured by the fact that if it is metaphysically impossible for one to know *Q*, then it is also impossible$_c$ for one to know *Q*. For the circumstantial notion of *possibility*$_c$ is also not supposed to quantify over worlds or scenarios that the notion of *metaphysical possibility* doesn't already quantify over.[102]

As is made clear by the considerations we just made, the fact that *Q* has false presuppositions at a world *w* is not the *only* fact that makes it impossible$_c$ for an inquirer to know *Q*, relative to *w*. Other factors that contribute to the impossibility$_c$ that the inquirer knows *Q* relative to *w* include more generally the unavailability of information that is enough to settle *Q* at *w* (as witnessed by the example of the reader's standing with respect to the question of *what I was thinking about on January 6, 1991 at 3 p.m.*), as well as the unavailability of the cognitive skills that they would need to have in order to be able to settle *Q* at *w*.

In other words, it can be impossible$_c$ for the inquirer to settle *Q* at *w* even though *Q* relies on no false presuppositions at *w*. The No False Presupposition Norm (NFP), then, does not entail the new Anti-Impossibility Norm (AIN). (AIN) is *strictly stronger* than (NFP).

Using the new models of inquiry (§4.3), we are able to correctly represent situations involving an inquirer and a question *Q* such that it is impossible$_c$ for that inquirer to know *Q* as follows. The inquirer's total body of information at world *w*, remember, is here captured by a set of possible worlds i_w—the set of worlds that are left open by her knowledge or evidence at world *w*. We saw that we can track the inquirer's learning new things through the income of new information by updates on i_w, namely, the elimination of some $v \in i_w$ by the newly acquired information. Information-

---

[100] See the Transcendental dialectic (division two) of the *Critique of Pure Season* ([1781, 1787] 1998).

[101] See for example Davis ([1958] 1982) on this, and also the reprint of Tarski ([1953] 2010).

[102] In other words: let *W* be the set of all possible worlds that the notion of *metaphysical possibility* quantifies over. Where *C* is the set of possible worlds that the (contextually restricted) notion of *possibility*$_c$ quantifies over, then, it must be the case that $C \subseteq W$.

gathering inquiry will chop i_w down into smaller sets. Here is one way, then, in which it can be impossible$_c$ for the inquirer to know Q: *there is no way* of updating i_w into a smaller set i^*_w such that $i^*_w \in Q$. In other words, no update in the inquirer's coarse-grained information in w will be such as to make it a complete answer to Q.

As we saw, the new models also feature a function r such that r_w is a set of representational vehicles that capture the *way* in which the inquirer represents world w as being (the world she is in). The inquirer can then make progress in armchair inquiry by reading off new vehicles out of her coarse-grained information i_w, and this will be tracked by updates in r_w, that is, by the transition from r_w to a new r^*_w that includes representational vehicles that were not included in r_w. Even though the inquirer's coarse-grained information may remain the same, again, her fine-grained information is thereby incremented—she comes to represent the world under new guises. So here is another way in which it can be impossible$_c$ for the inquirer to know Q (as presented under a guise g): *there is no way* of updating r_w into a bigger set r^*_w such that r^*_w contains a vehicle that constitutes a fine-grained answer to Q (as presented under g).

As it has become routine throughout our investigation, the latter point reminds us that the Anti-Impossibility Norm (AIN) should also be read as featuring a guise-parameter, just like the other norms. Its more explicit version is, then:

(AIN) One shouldn't inquire into Q *as presented under guise* g if it is impossible$_c$ for one to know Q *as presented under guise* g.

We have already found two motivations to include (AIN) in our system of norms of inquiry, namely, (i) it properly flags cases of misguided inquiry that the other norms are silent about (strength), and (ii) it entails others norms that we have already decided to include in our system (unifying power). We will now look at cases where it is not obvious that (i) applies.

§8.4 That the No False Presupposition Norm may or may not flag the impropriety of inquiry into indeterminate questions, and that yet another norm of inquiry may lie strictly between the Anti-Impossibility Norm and the No False Presupposition Norm

What about cases of indeterminacy, say, due to vagueness? Maybe some questions are *indeterminate*, in the sense that their maximal answers are neither true nor false (let us suppose for the moment that *there are* such questions). When a question is indeterminate in this sense, there is again no way in which the inquirer can come to know its answer.

Suppose it is indeterminate *whether Jack is bald at t*, on account of Jack being a borderline case of baldness at *t*. So neither of the two maximal answers to that question are true, namely, that *Jack is bald at t* and that *Jack is not bald at t*. Therefore, it is impossible$_c$ for any inquirer to know *whether Jack is bald at t*. For, since knowledge requires truth, it is impossible$_c$ for any such inquirer to know that *Jack is bald at t*, as well as it is impossible$_c$ for any such inquirer to know that *Jack is not bald at t*. So (AIN) will tell those inquirers not to look for the true answer to the question of *whether Jack is bald at t*. For there isn't a true answer to be found also in this case—we have here yet another example of inquiry that is bound not to achieve its own aim.

Does the No False Presupposition Norm (NFP) also issue a reproach to inquiry into such questions? The question of *whether Jack is bald at t* presupposes that *either Jack is bald at t or Jack is not bald at t*. That is because each of the two maximal answers to that question entails that *either Jack is bald at t or Jack is not bald at t*. In order for (NFP) to issue a reproach on that count, then, the target presupposition would have to be *false* (note incidentally that, if truth-value gaps are allowed, then (NFP) does not entail the Anti-Dissonance Norm (ADN) anymore). And whether that presupposition is false or not depends on what compositional principles of indeterminacy are the right ones.

According to most theories on the topic, the presupposition that *either Jack is bald at t or Jack is not bald at t* is not false. According to the most popular three-valued semantics—which feature an *indeterminate* value in addition to *true* and *false*—a disjunction of disjuncts whose values are indeterminate is itself assigned the indeterminate value.[103] So, in our example from above, the target presupposition would be *neither true nor false*, and therefore *not false*. Since that presupposition is not false, (NFP) issues no reproach to someone who inquires into *whether Jack is bald at t*, and that would be another example illustrating how (AIN) is strictly stronger than (NFP).

On an alternative, supervaluationist approach, the presupposition that *either Jack is bald at t or Jack is not bald at t* will even be deemed *true*, though its disjuncts are neither true nor false. The presupposition is deemed true here because, in all ways of sharpening the language, the relevant disjunction comes out true.[104] So here again

---

[103] See for example Kleene (1952), Łukasiewicz (1967) and Tye (1994).

[104] See Fine (1975), and also Keefe (2003, Ch. 7) on supervaluationist approaches to vagueness more generally. Though note that Fine is more directly occupied with the truth-value/lack thereof of *sentences*. We could represent indeterminacy in our framework of coarse-grained information by having the proposition that *Jack is bald at t* correspond to a bit of coarse-grained information understood as a set {*y, n*}, where *y* is the set of worlds where Jack is bald at *t* and *n* is the set of worlds where Jack is not bald at *t*—and the union of *y* and *n* may not equal the total space of possibilities *W*. Then we could talk of the sharpening of *Jack is bald at t* (instead of the sharpening of 'bald'), in terms of a set {y^s, n^s} where $y \subseteq y^s$ and $n \subseteq n^s$ even though this time $y^s \cup n^s = W$.

(NFP) would issue no reproach to someone who inquires into that question in world *w*, even though (AIN) would. (NFP) could only issue a reproach in such cases if a disjunction of disjuncts that are neither true nor false were itself to be false.

But here is another norm that *does* find fault with inquirers who inquire into such indeterminate questions, just like (AIN) does. We call it the No Untrue Answers Norm:

(NUA) One shouldn't inquire into *Q* if none of the complete answers to *Q* are true.

If some questions are indeed indeterminate, then (NUA) is strictly stronger than (NFP). In inquiring into the question of *whether Jack is bald at t*, for example, the inquirer is violating (NUA) without violating (NFP). For we have assumed that *Jack is bald at t* is neither true nor false, and therefore *Jack is not bald at t* is neither true nor false (those are the two maximal answers). But the target question doesn't rely on any *false* presupposition. So even if the impropriety of inquiry in these cases is not caught by (NFP), it is definitely caught by (NUA).

The good thing, however, is that the Anti-Impossibility Norm (AIN) *also* entails the No Untrue Answers Norm (NUA), and this is so regardless of whether some questions are indeterminate. For the complete absence of true answers to *Q* at a world *w* entails the unavailability of information that is enough to settle *Q* at *w*—and that again makes it impossible$_c$ for an inquirer to know *Q* at *w* (see the previous section).

In endorsing (AIN), then, we get both (NUA) and (NFP) for free. And that seems to be the best way to get both (NUA) and (NFP) into our system of norms, namely, through a norm that is strictly stronger than both of them. For neither (NUA) nor (NFP) are by themselves enough to issue a reproach into questions whose true answers are unknowable to the inquirer, even though those questions do not rely on any false presuppositions and they *do* admit of true answers (unlike indeterminate questions)

§8.5 Where we address two challenges to the inclusion of the Anti-Impossibility Norm into our system of norms of inquiry

In §6.3 we became concerned with the internal consistency of our system of norms of inquiry. In particular, we worried about the consistency between the Go Gather (GG) and the Go Figure (GF) norms, on the one hand, and norms such as the Ignorance Norm of Inquiry (INI), on the other. The former two norms are *proactive* norms, in the sense that they direct us *to inquire* into certain questions when we satisfy certain conditions with respect to those questions. And (INI) is a *prohibitive* norm, in the sense that it directs us *not* to inquire into certain questions when we satisfy certain conditions with respect to those questions.

The newly added Anti-Impossibility Norm (AIN) is also a prohibitive norm, as are its corollaries the No Untrue Answers (NUA) and the No False Presuppositions (NFP) norms. And one could try to argue against the consistency of the latter batch of norms with (GG) and (GF) as follows.

Suppose you want to settle *Q* and the information you have isn't yet enough to settle *Q*. So the Go Gather Norm (GG) tells you to engage in information-gathering inquiry, to collect more data. But suppose that, unbeknownst to you, it is impossible for you to know *Q* (say, because *Q* relies on a false presupposition). So (AIN) tells you not to inquire into that question, and that includes information-gathering inquiry. So you should inquire into *Q*, as per (GG), and you should not inquire into *Q*, as per (AIN). Our system simultaneously enjoins you to do and to not do something. Under the natural assumption that *you should not inquire into Q* entails that *it is not the case that you should inquire into Q*, furthermore, the system itself becomes inconsistent—for it says that *you should inquire into Q* while saying at the same time that *it is not the case that you should inquire into Q*. Therefore, one of these purported norms of inquiry must be false, and any system of norms comprising both (AIN) and (GG) is inconsistent.

We have seen a similar objection to (GG) before, in §5.4. It went as follows. A biographer wants to settle *what Marie Skłodowska Curie's favorite dish was*, but he doesn't have information that is enough to settle that question. And there are no entries on the topic in Marie's notebooks, or in the letters she sent to other people—no records whatsoever that our biographer could consult through which he could come to learn *what Marie's favorite dish was* (another gap in history that won't ever be filled). The antecedent of the conditional embedded in (GG) is true of our inquirer: there is a *Q* such that he wants to settle *Q* but he doesn't have information that is enough to settle *Q*. But it doesn't seem that he should look for information that will settle the issue, simply because *there isn't* such information for him to acquire. Why should one try to do the *impossible?* One shouldn't. But (GG) says one should, so (GG) must be wrong.

And our response to that objection was that (GG) is a wide-scope norm and, as such, our inquirer can come to satisfy it by ceasing to want to know *what Marie Skłodowska Curie's favorite dish was.* He doesn't have to attempt the impossible in order to abide by that norm. Assuming that he does is to commit the fallacy of 'factual

detachment'—that is, of drawing the conclusion that *It should be that q* on the basis of the normative premise that *It should be that* ($p \supset q$) and the factual premise that *p*.[105]

The objection to the consistency of our system from above commits the same fallacy. Just from the fact that you want to settle *Q* and that your information isn't enough to settle *Q*, given the wide-scope norm (GG), it doesn't yet follow that you should inquire into *Q*. We are not entitled to conclude, as suggested, that you should inquire into *Q* and you should not inquire into *Q*. In fact, there is a perfectly good way here for you to abide by *both* (GG) and (AIN) simultaneously here: you neither inquire into *Q* nor want to settle *Q*.

(Ideally, our system of norms will be such that the things you should do *all things considered* coincide with the things you must do in order to abide by *all* the wide-scope norms of the system at the same time. In order for this not to always be achieved through inaction and lack of wanting, however, there must be some conditions such that (i) you satisfy them, and it is not up to you to satisfy them, (ii) it is not the case that you shouldn't satisfy them, and (iii) they are also such that their satisfaction requires you to want to settle *Q*, for certain *Q*.)

Another consideration that seems to speak against the inclusion of (AIN) into our system of norms is the following. Sometimes it seems rational—even from an instrumental point of view—for one to inquire into a question *Q* though it is impossible$_c$ for one to know *Q*, maybe even knowingly so. For there are cases where all one is really aiming at is something like a well-informed high credence toward one of the maximal answers to *Q*.

To take an example from Falbo (forthcoming), consider an agent who is inquiring into *whether it will rain tomorrow*. By consulting the relevant data (the sky is clear, etc.), the agent becomes more confident than not that *it will not rain tomorrow*—though neither he nor even the meteorologists are in a position to know that *it will not rain tomorrow* (let us just assume that this is so). Our inquirer even acknowledges that he is not in that position, that it would be unrealistic for him to expect to know *whether it will rain tomorrow*. But that is okay, for all he wants is to make a good bet, so to speak. He inquires into that question and terminates his inquiry with a high credence that *it will rain tomorrow*, henceforth acting and making decisions on the basis of it.

---

[105] See Greenspan (1975) for the distinction between factual and deontic detachment. In general, an operator *X* allows for factual detachment when *X*(*if p then q*) and *p* entail *Xq*. An operator *X* allows for deontic detachment when *X*(*if p then q*) and *Xp* entail *Xq*. For many operators, factual detachment is found to be implausible. For example, in the literature about rational requirements, or requirements of coherence, both Brunero (2010) and Way (2011) defend the view that those requirements take wide-scope, without allowing for factual detachment (otherwise, wide-scope requirements would be objectionable in the same way that narrow-scope requirements are).

Cases such as this one seem to constitute a problem not only for (AIN), of course, but more fundamentally for the view on which we are grounding that norm, namely, that inquiry into *Q* is aimed at knowing *Q*.[106]

Our response, which shouldn't be surprising by now, is to challenge the very description of the case. It feeds off of the underdetermination or open texture regarding which questions the inquirer is more exactly inquiring into, given his behavior, and it assumes that this is ultimately left for the overall best theory to decide (best along dimensions such as explanatory power, deductive strength, simplicity, internal consistency). Where the objector has the subject inquire into *whether it will rain tomorrow*, we have him inquire into *whether it is likely that it will rain tomorrow*, or *whether rain is more likely than no-rain tomorrow*, or something along these lines. Where our inquirer cannot realistically expect to have categorical knowledge, he can more realistically expect to have probabilistic knowledge.[107]

We stand fast, then, to the Anti-Impossibility Norm.

---

[106] In fact, Falbo uses the example to object to the claim that inquiry aims at knowledge.

[107] Probabilistic knowledge might be construed either as knowledge of propositions about probabilities, or rather as knowledge whose content is *probabilistic content* (a set of probability spaces), as proposed by Moss (2018).

Chapter 9

§9.1 That sometimes it is instrumentally beneficial for inquirers to violate the norms of our system, though that is compatible with their being true norms of inquiry

Now it is time for us to address a more general concern about the system of norms that we have been building. It is not directed at this or that norm in particular, but to more than one of them at the same time.

The worry can be put as follows. We have endorsed norms that tell inquirers not to inquire into questions that have certain properties—for example, questions whose true answers are already known by them, or questions whose true answers are impossible for them to know. The target norms were advanced as *instrumental* norms of inquiry. They hinge on the possibility of achieving the constitutive goal of inquiry, on how it can be achieved, and on whether it has already been achieved. And it turns out that, sometimes, when we inquire into questions that have the target properties, we *do* further the goals of inquiry. We obtain inquiry-related benefits as a result of doing what is prohibited by those norms. But then it looks like they are not norms of inquiry at all. (A similar worry was addressed in §7.4, though it explored other goals that agents have *in general*, as opposed to their inquiry-related goals more specifically).

For example, the No False Presupposition Norm (NFP) says that one shouldn't inquire into the question of *which prime number is the largest one*. But what if inquiring into that question is the very way in which one comes to learn that *there isn't a largest prime number*? That surely looks like a benefit of inquiring into that question. The question is defective, to be sure, in that it relies on a false presupposition—but inquiring into it is the means through which the inquirer comes to realize that it is defective. And that is exactly the kind of epistemic benefit that inquirers aim at *qua* inquirers. In inquiring into that defective question, the agent thereby acquires new knowledge about numbers, after all. So why must we deem his activity of inquiry into that question counterproductive?

The contention presupposes that the true instrumental norms of inquiry must be all encompassing or exception-less, in the sense that it must never the case that an inquirer does something that is conducive to the overall goals of inquiry by violating the true norms of inquiry.

That presupposition, however, is ill-founded. Consider instrumental norms more broadly. One should take the necessary means to one's goals. But that is compatible

with the fact that, sometimes, by not taking the necessary means to one's goals, one ends up getting *other things* that one wanted to have (or they are simply welcome or desirable consequences to one). Suppose you didn't book a flight you were meaning to book. You have thereby saved some money. Plus, now you will be able to spend more time with your friends in town, etc. And these are good things. But it is still true that you ought to book the flight if you want to travel (give wide-scope to 'ought' here, too).[108]

There are many ways in which the violation of the norms of instrumental rationality can be beneficial to an agent. And yet they are still norms of instrumental rationality. In our case, there are many ways in which the violation of the norms of inquiry can be beneficial to the agent—'beneficial' not only in the very broad sense of the term, but also relative to the very goals that an inquirer has *qua* inquirer. Inquirers can come to know things, for example, as a result of violating those norms. They can maximize their overall goals *qua* inquirers even though they are acting in a counterproductive manner relative to the more specific goals of their inquiry into a particular question *Q*.

In the case introduced above, the agent benefits from inquiring into the question of *which prime number is the largest one*, in that he thereby comes to learn that *there isn't a largest prime number*. (NFP) says he is doing something that he shouldn't be doing, in the instrumental sense of 'should'. And the contention was: Why must that activity of inquiry be labelled counterproductive, seeing as it fosters the overall goals of inquiry for that agent, such as that of learning new things?

The answer is that his act of inquiry is counterproductive *relative to the goal of knowing which prime number is the largest one*. *That* is what the agent was trying to accomplish in inquiring into the target question. And he didn't manage to accomplish *that*—he was bound not to. The fact that he has thereby learned that *there isn't a largest prime number* doesn't make his activity of inquiry successful, not relative to the very goal that he had in performing that activity. That good consequence would only have counted as an *achievement* of her goal-directed activity if what he had been trying to do was to settle the question of *whether there is a largest prime number*. By assumption, however, that is *not* what he was trying to do. For we did not assume that he was inquiring into the question of *whether there is a largest prime number*. We have assumed, rather, that he was inquiring into the question of *which prime number is the largest one*. (NFP) issues a reproach to the latter act of inquiry, not to the former.

Additionally, let us also be more encompassing in our description of the inquiry-related benefits enjoyed by our inquirer in this example. One benefit of his activity of inquiry was, again, that he finally came to learn that *there isn't a largest prime number*.

---

[108] See Way (2010) for a defense of the wide-scope reading of the means-ends requirement.

But there is more to it than that. For now he is also in a position to avoid making the same mistake he was making before, that is, the mistake of searching for the true answer to the question of *which prime number is the largest one*. He can easily deduce that *there isn't such an answer* after all, and prevent both himself and other inquirers from using up their time and resources looking for what does not exist.

That he is now in a position to realize the error of his previous ways—as well as the errors of other inquirers' ways—is what we would expect to be the case were a norm such as (NFP) to be true. For not only does (NFP) entail that he was indeed in error, it can also be read off of our assessments or appraisals of each other's inquiring activities, as well as our assessments or appraisals of our own past inquiring activities.

Just as the present contention generalizes to the other norms of our system, so does the response. For example, one might object to the Go Gather Norm (GG) on the basis of the observation that, by never gathering more information that bears on a question *Q*, even though the inquirer wants to settle *Q* and doesn't yet have information that is enough to settle *Q*, the inquirer can dedicate his time and attention to *other* questions of his interest and thereby settle them. No doubt he is in this way furthering his overall goals of inquiry. But, no matter how well he performs elsewhere, this one thing is off about him, namely, he is not taking the means to get this one thing that he wants (he wants to know *Q*).

§9.2 That, relatedly, our norms are grounded on facts about the detraction and accretion of instrumental value to inquiry, relative to its own constitutive goal

Our response to the worry raised in the previous section can still be complemented as follows.

Consider prohibitive norms such as the Ignorance Norm of Inquiry (INI). An inquirer should not inquire into *Q* if she already knows *Q* (as presented under the same guise), and that is because her already knowing *Q* *deprives* the activity of inquiry into *Q* of instrumental value for her, relative to the very goal that is constitutive of that activity. The 'should' of (INI) hinges on that. The activity of inquiry won't promote *that* goal—namely, the goal of settling *Q*—when the inquirer already knows *Q*. But does not mean that such an activity is absolutely deprived of instrumental value for her, when we consider other goals the inquirer might have.

Or consider the Anti-Impossibility Norm (AIN). What makes it the case that an inquirer should avoid inquiring into a question *Q* such that it is impossible for her to know *Q* is that the impossibility of her knowing *Q* deprives the activity of inquiry into *Q* for her, relative to the very constitutive goal of that activity. But that again doesn't mean that such an activity is deprived of any instrumental value for her.

The instrumental value of an activity is a function of how useful that activity is as a means to the agent's goals. The same activity can have different instrumental values for different agents, even though they have the same goals. If both of two agents want to achieve some goal *x*, for example, but one of them has *already* achieved goal *x* (perhaps unbeknownst to her), then an activity that is supposed to promote goal *x* will have no instrumental value for the later agent, though it might have some instrumental value for the other one, as far as promoting goal *x* alone is concerned. Something similar happens with respect to two agents such that it is impossible for one of them to achieve *x*, but not for the other.

The fact that an agent has already achieved *x*, or that it is impossible for her to achieve *x*, detracts instrumental value of any activity that is supposed to further *x* for that agent—and it completely deprives any such activity of instrumental value, relative to goal *x* alone. That does not entail, however, that any such activity is completely deprived of instrumental value for that agent, for she might have *other* goals such that those activities are conducive to them. Accordingly, our instrumental 'should' is not an all-things-considered 'should' (if even there is such a thing).[109] It hinges only on facts regarding the possibility, realization and means to the *constitutive* goals of activities.

You *should not* inquire into *Q* when you already know *Q*, or when it is impossible for you to know *Q*, we say, because the activity of inquiry into *Q* won't be able to promote its own constitutive goal of settling *Q* under those conditions (in the first case because *Q* is already settled, in the second one because it is impossible for that to happen). But that doesn't mean that, *all things considered*, you should not inquire into *Q* when you already know *Q*, or when it is impossible for you to know *Q*. For the activity of inquiry into *Q* may still have some other instrumental value for you, seeing as it may promote other goals that you have (other than the constitutive goal of settling *Q* which is constitutive of that activity).

The 'should' of our proactive norms (GG) and (GF) is not an all-things-considered 'should', either. Those norms are true in virtue of the fact that satisfaction of the antecedent of the conditionals embedded in them, for some question *Q*, makes the activity of inquiring into *Q* valuable, even essential for the inquirer, relative to the very constitutive goal of that activity (the goal of the activity is the *same thing* that the inquirer is said to want to bring about in the antecedent). In the case of (GG): the activity of gathering new information of a specific sort (an activity that is aimed at settling *Q*) becomes essentially valuable in situations where the inquirer wants to know *Q*, but doesn't yet have information that is enough to settle *Q*. An analogous

---

[109] Relatedly, for some discussion and grounds for skepticism about all-things-considered 'ought', see Foot (1972), Tiffany (2007), Baker (2018).

explanation holds for (GF). These norms are not made false by the fact that information-gathering or armchair inquiry are not the *most valuable things* for inquirers to do in situations where they want to know the answers to their questions. By not inquiring into their questions they may be furthering other goals they have, other than the goal of settling their questions.

Such is the nature of the instrumental 'should' of our norms of inquiry. With these further clarifications at hand, we see once again that the worries raised in the previous section pose no threat to our norms. The fact that it is impossible for an inquirer to know *which prime number is the largest one* deprives the activity of inquiry into that question of instrumental value, relative to its very constitutive goal—even though that activity may promote some of her other inquiry-related goals.

Detraction and accretion of instrumental value from/to inquiry is not to be confused with negative and positive judgments that the inquirer herself makes about her own activity of inquiry. Instrumental value may be accreted to the activity of inquiry into *Q* for an inquirer who wants to know *Q*, even though she doesn't know or isn't sensitive to the fact that that activity has instrumental value to her. For example, her information may not be enough to settle *Q* without her knowing this, and so she is not in a position to appreciate how the activity of gathering new information has instrumental value for her.

Similarly, instrumental value may be detracted from the activity of inquiry into *Q* for an inquirer, even though she doesn't know or isn't sensitive to that fact. For example, it may be impossible for her to know *Q* without her knowing that it is, and so she is not in a position to appreciate how the activity of inquiry into *Q* isn't able to promote its own goal, thereby losing instrumental value relative to that goal.

§9.3 That our prohibitive and proactive norms entail their epistemically qualified counterparts

The last point from the previous section reminds us to distinguish the norms we have been advancing so far from what we could call their 'epistemically qualified' counterparts. For example, we are to distinguish:

(AIN) One shouldn't inquire into *Q* if it is impossible for one to know *Q*,

from:

(K-AIM) One shouldn't inquire into *Q* if one knows that *it is impossible for one to know Q.*

And we are to distinguish:

(INI) One shouldn't inquire into *Q* if one already knows *Q*,

from:

(K-INI) One shouldn't inquire into *Q* if one knows that *one already knows Q*.

And similarly for all the other prohibitive norms entailed by (AIN) and (INI) (we now omit the guise-parameter from the formulation of these norms to avoid clutter).

But we surely want to include (K-AIM) and (K-INI) in our system of norms of inquiry, too. To take the example from above, once the inquirer comes to know that *it is impossible for him to know what the largest prime number is*, he will also be in a position to recognize that his activity of inquiry into that question cannot promote its own aim. That activity will then be seen as a counterproductive activity, but this time also *from the inquirer's own perspective*, and not only ours'. Inquiry shouldn't be performed under such transparently hopeless circumstances.

But the good thing is that the system of norms we have been building already includes (K-AIM) and (K-INI) as theorems or consequences, though this hasn't been made explicit yet—and similarly with regard to all the other epistemically qualified versions of our prohibitive norms.

That is because (AIN) entails (K-AIM), and (INI) entails (K-INI), etc. It is easy to see why. In general, again, knowledge that *p* entails that *p*. In particular, that one knows that *it is impossible for one to know Q* entails that it is impossible for one to know *Q*. So if one shouldn't be such that (one inquires into *Q* even though it is impossible for one to know *Q*), then one shouldn't be such that (one inquires into *Q* even though one knows that *it is impossible for one to know Q*). For if one is such that (one inquires into *Q* even though one knows that *it is impossible for one to know Q*), then one is also such that (one inquires into *Q* even though it is impossible for one to know *Q*). If the latter is banned, so is the former. The same holds for the other prohibitive norms and their epistemically qualified counterparts.

What about our proactive norms, the Go Gather Norm (GG) and the Go Figure Norm (GF)? They go as follows, again:

(GG) If one wants to settle *Q*, but one's coarse-grained information is not yet enough to settle *Q*, then one should see to it that one gathers new information that bears on *Q*.

(GF) If one wants to settle *Q*, and one's coarse-grained information *i* is enough to settle *Q*, then one should see to it that one competently deduces some fine-grained information *x* such that $v(x) = a$, where *a* is the maximal complete answer to *Q* such that $i \subseteq a$.

These norms will have more epistemically qualified counterparts than the prohibitive norms. For example, one of the epistemically qualified counterparts of (GG) arises from the substitution of 'one knows that *one's coarse-grained information is not enough to settle Q*' for 'one's coarse-grained information is not enough to settle *Q*' in (GG). Or we can substitute 'one knows that *one wants to settle Q*' for 'one wants to settle *Q*' in it, too, which gives us yet another epistemically qualified version of (GG). Alternatively, we can make *both* substitutions at the same time—and that will give us yet another epistemically qualified norm of inquiry.

No matter which of these epistemically qualified versions of (GG) we formulate, however, they follow from (GG). The latter has the form $S(\varphi \wedge \psi \supset \sigma)$. Now for any χ such that χ entails that φ, (GG) will entail the norm $S(\chi \wedge \psi \supset \sigma)$. In particular, we can think of χ as an ascription of knowledge that φ, which entails that φ—if one knows that *one wants to settle Q,* then one does want to settle *Q*. Similarly, for any χ such that χ entails ψ, (GG) will entail the norm $S(\varphi \wedge \chi \supset \sigma)$. In particular, we can think of χ as an ascription of knowledge that ψ, which entails that ψ—if one knows that *one's information is not enough to settle Q*, then one's information is not enough to settle *Q*. And so on. Analogous points apply to (GF).

Some philosophers will perhaps choose to endorse these epistemically qualified counterparts without endorsing the norms that they epistemically qualify. Consider the Anti-Impossibility Norm (AIN) again. It says you shouldn't inquire into a question whose true maximal answer is unknowable to you. But what if you don't know that *it is unknowable to you*? Aren't you then normatively in the clear to inquire into that question? And, if that is the case, then (AIN) is false, even though (K-AIM) isn't.

But the reply is, again, that it *is* instrumentally ineffective for you to inquire into the target question—you are never going to settle it, no matter how the future turns out to be. You might have an excuse for inquiring into it, even though you shouldn't inquire into it, in the instrumental sense of 'should'.[110] You are *not* normatively in the clear to inquire into it, though perhaps you appear to yourself to be normatively in the clear to inquire into it.

This becomes clearer when we consider the position of an *assessor* who is looking at your situation from a more advantageous epistemic position. If *she* knows

---

[110] Relatedly, on the distinction between justification and excuses in epistemology, see the discussion by Williamson (forthcoming) and Greco (forthcoming).

that the true answer to your question is unknowable *to you*, then she is in a position to correctly judge that your inquiry into that question is bound to fail, that your activity of inquiry won't ever deliver the result you are thereby aiming at (settling your question). So she can perfectly describe you as making a mistake, or as being misguided, just as (AIN) would have it. Her assessment isn't affected by the fact that *you* don't know that the true answer to your question is unknowable to you. Suffice it that *she* knows that.

And even your future self, when in a better epistemic position, can make similar judgments about your past inquiring self. You look back and you realize that you were in the wrong path, inquiring into a question you could not possibly settle (a wrong path for you here is a path that does not take you where you wanted to be). Of course, both your future self and your present assessor may acknowledge that, quite understandably, this was the path that you saw fit to pursue back then. Your performance wasn't *blameworthy*, for you had no idea that it was the wrong path. But the wrong path it was.

Similar replies hold in response to those who want to endorse (K-INI) without endorsing (INI) itself, or those who want to endorse some of the epistemically qualified counterparts of (GG) and (GF) without endorsing (GG) and (GF) themselves.

§9.4 That inquirers are at least uncertain about the questions they are inquiring into

We have assumed that an agent who inquires into Q is driven towards knowing *Q*—she wants to know *Q*, or she has some kind of need to know *Q* (we have used the verb 'know' to precisify the verb 'settle' here). It is *because* the agent is so driven that she inquires, and it is constitutive of her activity that she satisfies that condition.

In order to be so driven, the agent need not explicitly think of herself *as knowing Q* in some possible, desired scenario—and certainly not through any particular linguistic vehicle involving the verb 'knows' or its cognates. It is rather that the agent wants or needs to bring about a certain state-of-affairs, which happens to be a state-of-affairs where she knows *Q*. That is the kind of possible world that she wants or needs the actual world to be. But such a state of affairs need not be described by that very agent as one where she knows *Q*. She might represent the desired or sought-after state-of-affairs in some other way, other than through linguistic vehicles. For example, Fido the dog from §1.1, who we countenanced as inquiring into *where the bone is*, may represent the state-of-affairs that it wants to bring about through mental imagery (a

state-of-affairs where it is in front of the bone, as seen from its own perspective, *ergo* a state-of-affairs where it knows *where the bone is*).[111]

Besides that volitional aspect of the activity of inquiry, we can also ask whether there are some other intensional conditions that an agent needs to satisfy (other than having the relevant kind of want or need) in order to count as inquiring into a question. We may be interested more specifically in whether there are *doxastic* conditions that an agent needs to satisfy in order to count as inquiring into a question, or doxastic conditions that are also constitutive of the activity of inquiry.

A natural suggestion is that inquirers always hold some kind of *interrogative attitude* towards the questions they are inquiring into.[112] Examples of such interrogative attitudes include the attitudes of wondering *how life began* and that of suspending judgment about *whether God exists*, for example. Notice how the verbs through which we ascribe such attitudes take interrogatives for their complements—hence the name 'interrogative attitudes'.

Often the inquirer inquires because she is in doubt about *whether such-and-such is the case*, or *when this-and-that happened*, etc. The attitude of *being in doubt* is also an interrogative attitude (contrast to the attitude of doubting *that* something is the case).[113] One can inquire into *what truth is*, for example, because one is in doubt about *what truth is.* From that inquirer's perspective, truth might be correspondence of language with reality, say, but it might also be a matter of coherence between representations. Similarly, one can inquire into *who bribed the policeman* because one is in doubt about *who bribed the policeman*. From that inquirer's perspective, Ana might have bribed the policeman, say, but Boris might have done it, too.

Here, the state of doubt teams up with the desire or drive to know the answer to a question and puts the agent in inquiring motion. In the first example, the inquirer is led to do philosophy. In the second one, the inquirer is led to perform a criminal investigation. Their topical differences aside, the state of being in doubt underlies both inquiring performances.

Is every inquirer in doubt about which of the answers to her question is true? It might be suggested that sometimes the inquirer is *almost certain* that one of those answers is the true one—maybe she even goes so far as believe it—though that doesn't constitute a state of being in doubt as to which of the question's answers is the true one. The inquirer heavily favors one of those answers, but she is trying to

---

[111] See Dror et. al (2022) for an empirical study on dogs' multi-modal mental imagery when they search for objects.

[112] The notion of an *interrogative attitude* here comes from Friedman (2013, 2017, 2019).

[113] See also Howard-Snyder (2013) and Rosa (forthcoming) for this contrast.

eliminate some lingering level of uncertainty that she still has about it (say, because it is very important to her that she doesn't make a mistake on the matter).

A safer suggestion, then, is that every inquirer is at least *uncertain* regarding which of the answers to her question is the true one. For at least two maximal answers *a* and *b* to *Q* (the question inquired into), it is uncertain to the inquirer *whether a*, uncertain to her *whether b*. To use the example from above, after narrowing down the suspects to Ana and Boris, it is uncertain to the criminal investigator *whether Ana did it*, uncertain to her *whether Boris did it*. It is uncertain to her *who did it*.

The inquirer can heavily favor one of the maximal answers to her question while still having uncertainty about it. Uncertainty has a better chance to be the kind of doxastic condition that underlies *all* inquiry, then. One can be uncertain without so much as being on the fence as to which of the maximal answers to the question is the true one—without so much as being in doubt as to which of those answers is the true one. Uncertainty is more *general* than doubt (same with suspended judgment and similar conditions of being doxastically divided). Whenever one is in doubt, one is uncertain, though one can be uncertain without exactly being in doubt.

Can we represent the relevant states of uncertainty using probability functions? It would be easy to add such functions to the possible worlds model of inquiry from §3.1. In addition to representing the situation of an agent through a set of worlds i_w that are compatible with what he knows/what his evidence is (his total body of coarse-grained information) at *w*, we'd also deploy a probability function pr_w to represent his levels of confidence in each of his questions' answers at *w*. Where *a* is a maximal answer to his question, then, that would be captured by the value that pr_w assigns to the intersection of i_w and *a*. From there on, we would need some way of measuring uncertainty that takes such values as inputs.

But the reason why that wouldn't work should be more or less obvious by now. We saw that the posit of representational vehicles or guises is crucial to capture important aspects of the activity of inquiry. We need to countenance the inquirer as manipulating fine-grained information to make sense of some of her inquiring activities and discoveries, if only because such activities do not hinge on the inquirer's coarse-grained information i_w being updated (i_w may remain unaltered from their beginning to their successful end). With only a probability function that takes sets of possible worlds as input, we won't be able to discriminate the inquirer's degree of uncertainty toward a given content when presented under one guise from her degree of uncertainty toward that same content when presented under an alternative guise (the same coarse-grained information as carried by two different vehicles).

Representing the element of uncertainty with our models of inquiry, then, would be a more complicated affair than simply adding probability functions to them. Using

the tools from the new models of inquiry (§§4.3–4.4), we could fetch an alternative function (not a standard probability function) that takes *pairs* of sets of possible worlds and representational vehicles as input—that is, bits of fine-grained information—and whose outputs are numerical values. We would then use those values to measure the inquirer's uncertainty concerning a question as presented under a certain guise. But let us dispense with such formal explications for now, and operate with an un-explicated, pre-theoretical understanding of the notion of uncertainty.

Someone who inquires into *how old the Earth is*, it seems, is at least uncertain about *how old the Earth is*. If she were to act as if she was inquiring into *how old the Earth is*, but she had no uncertainty about the issue whatsoever, we would be inclined to say that she is only *pretending* to inquire into that question, without actually inquiring into it.

Taken from the grammar of ascriptions of uncertainty, uncertainty is also an interrogative attitude: it can be uncertain to one *whether* such-and-such is the case, *when* this-and-that happened, *who* did this-and-that, etc. Given that much, then, we adopt the idea that uncertainty is the kind of interrogative attitude that underlies all inquiry. To inquire is, among other things, to do certain things or to go through certain thought processes because one is uncertain.

Uncertainty *bothers*, at least when there is some interest on the part of the agent to settle the question she is uncertain about. Inquiry is an attempt to kill the bother, which will hopefully bring the inquirer to know the answer to her question. Inquiry is the offspring of a doxastic condition—an interrogative attitude, minimally uncertainty—and a volitional condition—the want or need to know something. Inquirers always inquire because they satisfy such doxastic and volitional conditions, respectively.

§9.5 That there is just one sense in which one can inquire into a question one already has the answer to

In §1.1 we briefly touched on the problem of whether, in order to count as inquiring into a question at all, the inquirer must not already have the answer to her question (the 'must' there was descriptive, not normative—it is the 'must' of necessary conditions). Our quick response at that section was that it is possible for an inquirer to inquire into a question she already has the answer to. But our possibility claim there exploited the ambiguity of the expressions featured in it (when we utter 'it is possible that *p*' because '*p*' comes out true under one disambiguation, false under others). Now we are in a position to tackle the issue in more detail and make the required disambiguations.

We address the problem in two steps. First, we establish two different senses of 'having the answer to a question' that can be extracted from our theoretical framework.

Second, we ask if having the answer to a question in each of those senses conflicts with the doxastic and volitional conditions of inquiry established in the previous section.

In §§4.1–4.2, we fleshed out truth-conditions for sentences of types ‘the inquirer has the coarse-grained information that *p*’ and ‘the inquirer has the fine-grained information that *p*’, respectively. The truth-conditions for ‘the inquirer has the coarse-grained information that *p*’ went as follows: ‘*p*’ is true at all possible worlds that are compatible with the inquirer’s fine-grained information, say, her knowledge or evidence. It is having information in that sense, remember, that allows an inquirer to settle some of her questions through sheer deductive reasoning, from the comfort of her armchair. She already has the answer to her question before even settling it herself.

The truth-conditions for ‘the inquirer has the fine-grained information that *p*’, however, went as follows: there is a representational vehicle *x* that is similar to ‘*p*’ under certain dimensions—that is, a vehicle *x* that belongs to the class *f*(‘*p*’)—and *x* is part of the inquirer’s representation of reality. Since any vehicle that belongs to *f*(‘*p*’) carries the coarse-grained information that *p* (or: for any $y \in f(\text{‘}p\text{’})$, $v(y) = v(\text{‘}p\text{’})$), the fact that the inquirer has the fine-grained information that *p* entails that she has the coarse-grained information that *p*. But since the fact that ‘*p*’ is true at all possible worlds compatible with the inquirer’s knowledge or evidence does not entail that there is a vehicle $x \in f(\text{‘}p\text{’})$ such that *x* is part of her representation of reality, the fact that the inquirer has the coarse-grained information that *p* does *not* entail that she has the fine-grained information that *p*.

Now we can use those very truth-conditions to disambiguate ‘the inquirer has the answer to *Q*’. Under one disambiguation of ‘has the answer to *Q*’, the target sentence comes out true just in case the inquirer has the coarse-grained information that *p*, and *p* is the true maximal answer to *Q*. Under an alternative disambiguation of ‘has the answer to *Q*’, the target sentence comes out true just in case the inquirer has the fine-grained information that *p*, which is the true maximal answer to *Q* presented under some guise $x \in f(\text{‘}p\text{’})$. Let us now explicitly import the disambiguation into our grammar and use the phrases ‘has the coarse-grained answer to *Q*’ and ‘has a fine-grained answer to *Q*’, respectively, to distinguish those two cases (the expressions ‘coarse-grained answer’ and ‘fine-grained answer’ were already put to good use throughout Chapter 6—and notice now that we use the definite article ‘the’ in the former case, but the indefinite article ‘a’ in the latter one, because there can always be more than one true maximal *fine*-grained answer to the same question).

We have now taken the first of the two steps mentioned above. Now on to the second one. Does having the *coarse-grained* answer to *Q* conflict with the volitional and doxastic conditions of inquiry into *Q*? And does having the *fine-grained* answer to

Q (as presented under some guise *g*) conflict with the volitional and doxastic conditions of inquiry into *Q* (as presented under that same guise *g*)?

It is clearly possible for one to want to know *Q* and be uncertain about *Q* when one already has the coarse-grained answer to *Q*. One has the answer to *Q*—but this is not yet explicit to one. For all it takes for one to have the coarse-grained answer to *Q* is that such an answer is true at all possible worlds that are left open by what one knows or what one's evidence is. Infinitely many propositions are true at all those worlds, but there isn't a one-to-one correspondence between such propositions and the representational vehicles that constitute the inquirer's representation of reality.

This is, again, the situation of the inquirer who is in a position to settle *Q* using armchair methods, even though she hasn't settled *Q* yet. One of the *reasons why* she is able to settle *Q* through armchair inquiry is that she already possesses the coarse-grained answer to her question. Her curiosity or her pull towards knowing *Q*, together with her uncertainty regarding *Q*, may eventually lead her to flesh out the true maximal answer to *Q* explicitly (under some guise or other), and she can do that by using nothing but her own thinking skills. She doesn't strictly need to acquire new coarse-grained information—she already has all the coarse-grained information she needs to settle her question.

We could also describe the situation as follows. Initially there is a *mismatch* between the inquirer's doxastic state and how her coarse-grained information bears on her question. The inquirer can even be completely uncertain (in doubt) as to which of the maximal answers to her question is the true one—and, yet, the question is *closed* relative to the coarse-grained information possessed by her. With enough cognitive skills, she can kill her uncertainty by adding new representational vehicles to her picture of how the world is like via competent reasoning, at which point her doxastic state will match the way in which her coarse-grained information bears on the question.

But that very explanation suggests a negative answer to the question of whether it is possible for one to have the *fine-grained* answer to Q, as presented under guise *g*, while at the same time being uncertain about *Q*, as presented under *g* again.

For the explanation suggests that, when the inquirer settles *Q* under guise *g*—in which case some declarative rendering of *g* is added to her representation of reality (§6.4)—her uncertainty about *Q* ceases to exist, or at least her uncertainty about *Q* *as presented under guise g* ceases to exist (the inquirer might still be uncertain about *Q* as presented under some other guise). Of course, the explanation is careful enough not to state that that is *always* the case—as witnessed by the use of 'can' in the last sentence from the previous paragraph. Still, it has us wonder whether, in having a fine-grained answer to *Q* as presented under *g*, an agent isn't thereby *prevented* from being uncertain about *Q* as presented under *g*. Perhaps it is psychologically impossible for

one to have a fine-grained answer to *Q* under *g* and be at the same time uncertain about *Q* under *g*.

Let <*x*, *p*> be a piece of fine-grained information possessed by the inquirer, and let it be a fine-grained answer to *Q*, as presented under *g*. So *x* is a representational vehicle, which we have described as a declarative rendering of *g*, such that it carries the coarse-grained information that *p*. The latter is a maximal answer to *Q*. The inquirer, then, has a fine-grained answer to *Q*, as presented under *g*. Given on our truth-conditions from above, that means that *x* is part of the inquirer's representation of reality. And the question now is whether this is compatible with the inquirer's being uncertain about *whether p* when *p* is presented under guise *x*.

For example, can you posses the fine-grained information that *this book was written by a human being*, as presented under the guise of the sentence just italicized, while it is at the same time uncertain to you *whether this book was written by a human being*, as presented under the guise of the interrogative just italicized? Supposing the former to be the case, it follows that some vehicle that is similar to the sentence 'this book was written by a human being' is part of your representation of how the world is like (similar under certain important dimensions, as we saw in §4.2). You represent the world as being that way—and you do it through that kind of linguistic representation. Using such representations, you form a picture of reality. You rely on your picture of reality to act and draw further inferences.

But then it would seem that, if you were to be uncertain as to *whether this book was written by a human being*, as presented under the interrogative just italicized, then you would *not* represent the world as being such that *this book was written by a human being*, as presented under the declarative sentence just italicized (though it could still be very probable from your perspective that *this book was written by a human being*, as presented under the declarative just italicized). Your uncertainty wouldn't allow the latter vehicle to be fully incorporated into your picture of reality, so that you could act and draw inferences on the basis of it. You would rather be *hesitating* to include that vehicle into your picture of reality. If you do represent the world as being such that *this book was written by a human being* under that very guise, then, it follows that is not uncertain to you *whether this book was written by a human being*, as presented under that very guise.

One objection that might be raised here is that, if the reasoning from the previous paragraph were sound, then you would have to be *certain* that *this book was written by a human being*, at least as presented under the guise 'this book was written by a human being', in order to represent the world as being that way under that guise.

But from the fact that one is *not uncertain* that *p* it does not follow that one is certain that *p*, even where *p* is presented under the very same guise in both cases. For

one can be neither certain nor uncertain that *p* (as presented under guise *x*), at least where a notion of *psychological* or *subjective* certainty is concerned.[114]

In particular, from the fact that one represents the world as being such that *p*, in that one's representation of reality includes a representational vehicle *x* with *v*(*x*) = *p*, it does not follow that one is certain that *p*, in the sense of being maximally confident that *p* is the case, even when presented under guise *x*. What follows is just that one is *not uncertain* that *p* is the case, when presented under guise *x*. One deploys *x* to make decisions and draw inferences—one *relies* on *x* in a way that someone who is uncertain that *p* when presented under guise *x* wouldn't.

It seems, then, that one cannot possess the fine-grained information that *p* while being uncertain about whether *p* under the very same guise. We conclude that there is only one clear sense in which one can have the answer to a question *Q* while being uncertain about *Q*, namely, the sense in which one possesses a *coarse-grained* answer to *Q*, while being uncertain about *Q*.

To go back to the question that started this section, then, we can now state the following conclusions. It is possible for one to inquire into *Q* when one already has the *coarse-grained* answer to *Q*. But it is not possible for one to inquire into *Q*, as presented under *g*, when one already has the *fine-grained* answer to *Q*, as presented under *g*. And that is because inquiry into *Q*, as presented under *g*, requires at least uncertainty about *Q*, as presented under *g.* But when one has the fine-grained answer to *Q*, as presented under *g*, one is not uncertain about *Q*, as presented under *g.* (Similar conclusions could be drawn about the volitional condition of inquiry).

---

[114] See Stanley (2008) for a more or less standard way of distinguishing between subjective or psychological certainty and epistemic certainty.

Chapter 10

§10.1 That there are also norms of inquiry based on the assessment of the attitudes that are constitutive of it, but that those norms are not instrumental norms of inquiry

As we saw, being in doubt or at least uncertain is part of the impetus of inquiry, together with some kind of need or desire to know (say, some kind of curiosity). It is *because* inquirers are in those kinds of states that they so much as inquire into certain questions.

Given their importance, then, is the epistemological assessment of the interrogative attitudes that underlie inquiry (doubt, suspension, minimally uncertainty) relevant to the assessment of the activity of inquiry itself?

It is clear that inquiry itself, understood as an act or activity, does not inherit all the kinds of appraisal or assessments that apply to the attitudes that are constitutive of it. Consider for example evidentialist notions of epistemic justification, as it applies to doxastic attitudes: a given doxastic attitude regarding *whether p* is justified relative to a total body of evidence *E* just in case that attitude matches or fits the way in which *E* bears on the question of *whether p*.[115] Using some such notion, we might say things like: since the inquirer's total evidence doesn't sufficiently support either of *p* and ¬*p*, she is *epistemically justified* in being in doubt as to whether *p* is the case. But her *act* of inquiring into *whether p* cannot be said to be epistemically justified. It is only doxastic attitudes that can be justified in that evidentialist sense, not acts or activities.

That is not to say, of course, that there is no sense in which inquiry can be said to be justified. Indeed, inquiry may be said to be *unjustified* in certain cases—in whatever sense it can in principle be said to be justified or unjustified—even though the interrogative attitude that underlies it is epistemically *justified*. That may hold, say, because the subject knows that it is impossible for her to settle her question (which makes the activity of inquiry unjustified for her), though her evidence doesn't support any of the maximal answers to her question more than it supports the others (which makes her interrogative attitude justified).

---

[115] An evidentialist account of epistemic justification is influentially put forward by Feldman and Conee (1985). See McCain (2014) for a more recent variation of the view.

Despite the lack of transmission of epistemic status from the interrogative attitudes that underlie inquiry to the act of inquiry itself, however, we might still want to take the following to be a true norm of inquiry—call it the *Norm of Inquiring Attitudes*:

(NIA) One should inquire into *Q* only if one's interrogative attitude towards *Q* is epistemically justified.

The interrogative attitude towards *Q* mentioned in the consequent of the conditional embedded in (NIA)—regardless of which attitude it more exactly is (doubt, or suspension of judgment, or just uncertainty)—is supposed to be the attitude that is constitutive of the activity of inquiry into *Q* that is mentioned in the antecedent.[116]

Suppose, for example, that an agent inquires into *whether some philosophers are feminists*, and that the interrogative attitude that underlies her activity of inquiry is an attitude of being in doubt as to *whether some philosophers are feminists*. Nevertheless, she knows that *Alex is a philosopher and a feminist*, and that *Mary is a philosopher and a feminist*. These facts, coupled perhaps with the fact that the agent is sensitive to the relevant logical relations between these contents, entail that the attitude of being in doubt as to *whether some philosophers are feminists* is not epistemically justified for her. It follows, then, that the agent violates (NIA). For she is inquiring into *whether some philosophers are feminists* while lacking justification to be in doubt about *whether some philosophers are feminists*.

Which terms of assessment could we use to assess activities of inquiry that violate (NIA)? Some options that suggest themselves are 'not proper', 'unwarranted', 'irrational' and 'blameworthy'. Accordingly, we might want to say that one's activity of inquiry counts as *proper*, or *warranted*, or *rational*, or *blameless*, only if it abides by the norm (NIA).

Compare that to the prohibitive norms that we have advanced before, crucially the Ignorance Norm of Inquiry (INI) and the Anti-Impossibility Norm (AIN). We call inquiry that violates those norms 'instrumentally counterproductive' or 'inefficacious'. We might *also* want to say that activities of inquiry that violate them are 'not proper'—but all we mean by that is that those activities are instrumentally counterproductive or inefficacious vis-a-vis the constitutive goal of settling the question inquired into (constitutive of that inquiry). In the case of a violation of (NIA), however, we mean something else when we say that the target activity of inquiry is 'not proper'. In flagging a violation of (NIA) as 'not proper' or 'improper', it is not as if we are assuming that

---

116 A more precise formulation of (NIA), then, is that one should be such that: if one inquires into *Q* and IA(*Q*) is the interrogative attitude that is constitutive of one's activity of inquiry into *Q*, then IA(*Q*) is epistemically justified.

inquiry whose underlying interrogative attitude fails to be justified is instrumentally counterproductive or inefficacious. Similar things could be said about 'unwarranted'.

'Irrational' can perhaps also be used to flag violations of the Ignorance Norm (INI), and we have indeed taken that norm to be a norm of instrumental *rationality*. But probably many will find it incorrect to label all activities of inquiry that violate our other norm, the Anti-Impossibility Norm (AIN), as 'irrational', even if they don't find it incorrect to label activities of inquiry that violate (INI) as 'irrational'. A theory of rationality that vindicates both of these labelings across the board will perhaps be perceived as an overly radical externalist theory of rationality. But it will in any case be a theory of *instrumental* rationality. When we call an activity of inquiry 'irrational' on account of the fact that it violates the Norm of Interrogative Attitudes (NIA), we do *not* thereby mean to judge that activity to be *instrumentally* irrational.

Finally, it also seems inappropriate to predicate 'blameworthy' of any inquirer who violates our norms (INI) and (AIN). That an agent inquires into a question whose true maximal answer is unknowable to him is not itself enough to license an ascription of blame to that inquirer (he may justifiably think that the answer *is* knowable to him). But 'blameworthy' applies more naturally to violations of the Norm of Inquiring Attitudes (NIA). An agent who inquires into *Q* without having justification to have an interrogative attitude toward *Q* can be judged blameworthy in a similar way that an agent who acts on the belief that *p* without having justification to believe that *p* is judged blameworthy. In both cases, we have an action based on a doxastic attitude that is unjustified.

The point is that there is substantial difference between the terms of appraisal that we use to flag violations of the Norm of Inquiring Attitudes (NIA), on the one hand, and those that we use to flag violations of our instrumental norms of inquiry, such as the Ignorance Norm of Inquiry (INI) and the Anti-Impossibility Norm (AIN).

Accordingly, the 'should' of (NIA) is not exactly the same as the 'should' of our instrumental norms of inquiry (AIN), (INI), (GG), (GF) and their consequences.[117] That the inquiring attitude towards *Q* is epistemically justified is not a means to the goal of inquiring into *Q*, and neither does its non-justification render inquiry into Q into a counterproductive activity. But such are the criteria that we have used to decide if candidate instrumental norms of inquiry are true. Those criteria are of no use to judge whether (NIA) is true. In fact, it isn't even very clear what could be the truth-maker of a norm such as (NIA)—though it is very clear what that is in the case of our instrumental norms (facts about the detraction and accretion of instrumental value to the activity of inquiry, as seen in §9.2).

---

[117] In saying this, we are not thereby committed to the claim that 'should' is *ambiguous*. It might rather be that it takes one modal basis from the context of assertion in one case, another in the other, in the sense of Kratzer (1977, 1981).

None of that speaks *against* our endorsement of (NIA), of course. But we have been occupying ourselves with the instrumental aspects of the normativity of inquiry here, and our norms are restricted to that aspect. The 'should' operator of (NIA) is of a less tractable and more obscure nature, and we have to leave it for future work to explore its connections to our instrumental 'should'. For now, it suffices to note that (NIA) is *not* the kind of norm that is supposed to be included in the system of norms of inquiry that we have been building throughout this investigation—which is a system of purely instrumental norms.

§10.2 Where a purported norm against inquiry into pseudo-questions is seen to be misguided

In the last session we have discussed a purported norm such that, even though it may also be a norm of inquiry, it is not an instrumental norm of inquiry. And our primary concern here was with building a unified system of instrumental norms of inquiry, the question of how the norms of such a system interacts with other kinds of norms being relegated to future investigation.

But there is yet another norm that looks similar to the ones we have been including in our system, namely, a norm that is violated whenever we occupy ourselves with *pseudo-questions*. A first attempt at formulating such a norm goes as follows (where 'NAPS' abbreviates 'Norm Against Inquiry Into Pseudo-Questions'):

(NAPS*) One shouldn't inquire into pseudo-questions.

We might try to justify (NAPS*) in roughly the same way we have justified the No False Presupposition Norm (NFP). A question with a false presupposition is a kind of *defective* question—for there isn't a true maximal answer to it. Where *Q* relies on a false presupposition, then, one will never achieve the constitutive goal of the activity of inquiring into *Q* (one will never settle *Q*, know *Q*). Similarly, a pseudo-question is a defective question—there isn't any particular maximal answer to it that is true, either. Where *Q* is a pseudo-question, then, one will never achieve the constitutive goal of the activity of inquiring into *Q*. Since that was enough to justify (NFP), it should also be enough to justify (NAPS*).

But that line of thought is at least partially misguided. A pseudo-question is not a defective question, for it is not even a question to begin with. Just like a pseudo-scientist is not a scientist, a pseudo-question is not a question (a pseudo-scientist only *pretends* to be a scientist, without actually being one). It is not even *possible* for one to inquire into a pseudo-question, for inquiry is necessarily inquiry into a question. Accordingly, (NAPS*) boils down to: one shouldn't do the impossible. But that is not

exactly what we are thinking of when we say that inquirers should avoid occupying themselves with pseudo-questions.

Why say that the attempted justification is only *partially* misguided? The reason is that there isn't indeed any particular true maximal answer to a pseudo-question, and no one will ever know the true maximal answer to a pseudo-question. But that is just because a pseudo-question is not a question, *ergo* there is no such a thing as the true maximal answer to it (just like there is no such a thing as the true maximal answer to a t-shirt, say).

If a pseudo-question is not a question, then what is it? Let us explore the following idea: a pseudo-question is a certain kind of *interrogative sentence* as uttered in a certain type of situation. Some interlocutors in the context of that situation will take it or act as if that interrogative expresses a particular question, thought it actually doesn't. We are thinking again of questions as the semantic contents of interrogatives (in our case, we are thinking of that content as the set of all complete coarse-grained answers to the question, or the downward closed set of its maximal answers). There is no guarantee, however, that an interrogative construction will express some particular question. When that doesn't happen, we call the interrogative construction a pseudo-question, and the act of uttering it the act of posing a pseudo-question.

There is more than one way in which an interrogative sentence may fail to express or convey a question in the context where it is uttered. The sentence might feature a singular term, say, that hasn't been assigned a referent—neither by the causal links and conventions that are responsible for assigning referents to the expressions of the target language, nor by the parameters of the context where it is used. To borrow and modify an example from Stalnaker (1970), suppose a friend asks you 'Is he a fool?' while pointing in the direction of Daniels and O'Leary, but without doing enough (or even intending) to determine whether 'he' refers to Daniels or to O'Leary. In this case, the target token of 'Is he a fool?' expresses neither the question of *whether Daniels is a fool* nor the question of *whether O'Leary is a fool*. If you were to search for the true answer to the question expressed by that instance of 'Is he a fool?', by reasoning from the information you already have, or by collecting more information (or both), you would be occupying yourself with a pseudo-question. And that would be another search in vain.

Pseudo-questions in the present sense may also fail to express questions because they deploy ambiguous expressions, and the situation where they are uttered does not allow any interlocutors to disambiguate them. Even the speakers who utter them may not know more specifically what they meant to ask. A perhaps not atypical situation of this kind is one where the speaker utters 'What is the meaning of life?',

without providing enough information to allow his interlocutors to disambiguate the expression ‘meaning’ in any particular way.

Something similar to that happens with the use of interrogatives that deploy technical terms—terms whose meanings have to be *postulated* within a specific linguistic community, instead of being borrowed from an already consolidated use in ordinary language. As an example from the area of epistemology, consider an utterance of ‘Does misleading higher-order evidence destroy epistemic justification?’, where the context and the conversational background of the utterance do less than specify how ‘epistemic justification’ is to be interpreted (it could mean fit with the total evidence, or reliability of belief-formation, etc.).

The legitimacy of the examples just offered aside, the idea is to describe the phenomenon of raising pseudo-questions as follows. Someone *tries* to raise a question by uttering an interrogative sentence *s*. But the context and conversational background of that utterance are not enough to determine which particular question is expressed by that uttered token of *s*, or the meaning-determining factors of the language deployed by the speaker do less than determine which particular question is expressed by it. There is more than one question Q_1 and Q_2 such that, under one disambiguation or contextual parameter-fixing or precisification, *s* expresses Q_1 and, under an alternative disambiguation or contextual parameter-fixing or precisification, *s* expresses Q_2.

Often it will be the case that the true maximal answer to Q_1 is *p* and the true maximal answer to Q_2 is *q*, where $p \neq q$. Even though *p* and *q* are two different answers, however, they are mutually consistent—otherwise they couldn’t both be true. But then the members of a purported dispute, which is elicited by the utterance of the target interrogative *s*, might wrongly think that they disagree with each other. For they utter declarative renderings of *s* whose grammar makes it look like they are contradicting each other. One of them utters ‘*p*’, because she thinks that *p* is the true maximal answer to Q_1. In response, the other one utters ‘¬*p*’, because she thinks that *q* is the true maximal answer to Q_2, and that the former speaker’s utterance of ‘*p*’ was meant to express an answer that is incompatible with *q* (an answer that entails ¬*q*). Both of them think that each other’s utterances convey incompatible answers. But they are actually offering answers to different questions, which can be expressed by the same vehicle, namely, the very interrogative *s* that elicited the debate between them. And their answers are actually consistent with each other.

So the posing of pseudo-questions is prone to generate ‘merely verbal’, non-substantial disputes. And that is because the meanings of the target interrogatives are underspecified by the situation of their use.

Philosophers attentive to the phenomenon might be led to think that many or most debates in philosophy are but examples of it. They are liable to characterize

philosophical disputes as merely verbal disputes, the questions of philosophy as pseudo-questions. They may claim that philosophical problems are resolvable (if at all) via settling issues about the meanings of words, or maybe by decisions about which conceptual scheme to adopt.[118] Some such diagnoses fly in the face of the fact that these mutual misunderstandings—the use of the same linguistic vehicle to convey different meanings—are pervasive across *all* areas of human inquiry, and that there are many problems in philosophy that are *not* merely verbal (hopefully, the problem of establishing what the instrumental norms of inquiry are is an example of a problem that is not merely verbal). Pseudo-questions and merely verbal disputes are not exclusive to philosophy, and neither is philosophy simply made out of them.

If pseudo-questions are (tokens of) interrogative sentences as they are used in a particular context or situation, then what is to be made of the norm that one shouldn't occupy oneself with pseudo-questions? What is it for one to 'occupy oneself' with pseudo-questions? We tackle these issues next.

§10.3 Where we explicate the vaguely formulated idea that one shouldn't occupy oneself with pseudo-questions

In the previous section, we saw that the norm that tells us to avoid pseudo-questions cannot be interpreted as saying that we shouldn't *inquire* into pseudo-questions. For pseudo-questions are not questions at all, and it is therefore not even in principle possible for one to inquire into pseudo-questions.

After establishing that much, we have started exploring the idea that pseudo-questions are tokens of *interrogatives sentences*, as uttered in some particular kind of situation. Grammatical constructions as they may be, the fail to express or convey any particular question. Different questions are different sets of answers. If the conversational background, the operative linguistic conventions, or the context of utterance were fixed in one way, the pseudo-question would have one set of answers for its content—and if those elements were fixed in some other way, it would have an alternative set of answers for its content. But those elements were neither fixed in this way nor in that way in the situation of utterance of the target interrogative, so that there is no such a thing as *the* true maximal answer to the question expressed by it. Such is the way in which we have characterized the phenomenon of an inquirer's raising pseudo-questions.

But now we have to find a way of explicating the roughly formulated norm to the effect that inquirers shouldn't occupy themselves with pseudo-questions, seeing as

---

[118] See, for example, the proposal advanced by Thomasson (2009) regarding questions of ontology.

(NAPS*) from above didn't quite do the job. To a first approximation, using the notion of pseudo-question just introduced, the norm would say that inquirers shouldn't occupy themselves with interrogatives of a certain sort. But even that is still a bit obscure. What would it be for an inquirer to 'occupy herself' with an interrogative sentence? For example, we wouldn't count the inquirer's *spelling* that sentence as occupying herself with it in the relevant sense. So we still have work to do here.

Before we move on to that, however, an important observation is in order. In characterizing the phenomenon of raising pseudo-questions above, we have described the inquirer as *uttering* the target interrogative construction *s*. But *s* might as well have been *written down*, or *tokened* in the inquirer's thought (his inner voice), instead of being said out loud, without making much difference to our characterization. And, presumably, an inquirer can occupy himself with a pseudo-question even though he is not in a dialogue situation—he is rather reading or engaging in soliloquy, or simply thinking for himself. We should broaden our characterization to include such phenomena as well: the meaning of the interrogative is underspecified by the linguistic conventions, or by the context or situation where it was *written down*/*tokened in the inquirer's thought*, etc. We leave these further details aside for now, however, and use the phenomenon of raising pseudo-questions via the utterance of an interrogative as a paradigm of the phenomenon, thus hoping that these further details won't affect our main points.

It might be suggested that something like the following will do as a more precise formulation of the norm that directs us not to occupy ourselves with pseudo-questions (where 'IN' stands for *Interrogation Norm*):

(IN) One shouldn't take oneself to be raising a question through an interrogative *s* when *s* doesn't really express any particular question.

The idea behind this norm is that one is doing something wrong when one *attempts to pose some specific question* through an interrogative that turns out to be empty of content (it is a pseudo-question). The interrogative *could* have expressed some question or other, of course, but the situation where the speaker deploys it does not satisfy the conditions that are needed for it to do that. To the extent that the speaker takes herself to be raising a question through that interrogative, he is making a mistake —that is what is wrong with his attempt.

Correct as that may sound, it doesn't seem like (IN) is the kind of norm that we have in mind when we say that inquirers shouldn't occupy themselves with pseudo-questions. Here, to *occupy oneself* with a pseudo-question is just to use it in an attempt to raise or pose a question. Where the attempt to raise or pose a question consists of

the *utterance* of an interrogative *s*, (IN) comes into play as a norm for *speech acts* of a certain kind, namely, speech acts of questioning or interrogating (if we were thinking of the target attempt as consisting of the tokening of *s* in thought, it would be a norm for a certain kind of *inner speech act*). But then (IN) isn't at all similar to the instrumental norms of inquiry we have been advancing here. One can perform such speech acts without so much as trying to find the true answer to a question, which is the kind of attempt that underlies all activity of inquiry (in contrast to the attempt of raising or posing a question).

To help make the point, notice that (IN) is analogous to a norm for the speech act of assertion, the *Assertion Norm*:

(AN) One shouldn't take oneself to be asserting a proposition through a declarative sentence *s* when *s* doesn't really express any particular proposition.

Just like there are interrogative sentences such that, depending on the situation of their use, they do not express any question, so there are declarative sentences such that, depending on the situation of their use, they do not express any proposition (or they do not carry information). The kinds of attempt that both (IN) and (AN) are used to assess are attempts to make meaningful use of language—attempts to use linguistic constructions in such a way as to allow them to *convey content*. We make such attempts in all sorts of situations: when we tell our friends what we did and how we feel, when we lecture, etc.

And, whereas such communicative attempts are relevant and sometimes even crucial to our inquiring pursuits, we meant to talk about something else when we used the phrase 'occupy oneself with a pseudo-question' above, other than the mere attempt to pose or raise a question via the utterance of an interrogative. What an inquirer is trying to do when he 'occupies himself with a pseudo-question' is more than just trying to use language meaningfully—in this case, interrogative sentences. We had something more intimately connected to inquiry in mind when we used that phrase.

A philosophy professor who tries to pose a question to his students in the classroom, for example, but who is actually (unknowingly) posing a pseudo-question to them, violates (IN) and is for that reason liable to criticism. He might be uttering the target interrogative for purely pedagogical purposes, however, without himself attempting to *settle* any question. He is simply trying to explain or show his students what a purported problem is. Of course, since there actually isn't any particular problem or question that is conveyed by his interrogative (a pseudo-question), he is definitely doing something wrong. But in making this assessment of his performance we are not assessing him negatively *qua* inquirer, aware as we are that he is not really trying to

settle any question. We are rather assessing him negatively *qua* instructor, and more basically *qua* communicator of content to an audience.

So (IN) doesn't quite get to the core kind of attempt that is to be assessed in light of the norms of inquiry, which is the attempt to find (to know) the true answer to some question. With that in mind, a alternative idea is that the norm that tells us to avoid occupying ourselves with pseudo-questions amounts to something with the following form (where 'NAPS' abbreviates 'Norm Against Our Occupying Ourselves With Pseudo-Questions'):

(NAPS) One shouldn't attempt to find the true answer to the question expressed by an interrogative *s* such that *Φ*(*s*).

Here, '*Φ*' is an open sentence that characterizes a pseudo-question, understood as an interrogative of a certain kind. '*Φ*' will presumably be a big disjunctive sentence, each disjunct of which ascribes linguistic/semantic properties and relations to objects that satisfy its free variable '*x*'. It may for example look like this: '*x* features indexical expressions, though the context where *x* is uttered doesn't provide the elements that are needed to determine their referents, or some expressions of *x* are ambiguous, though the conversational background where it is uttered doesn't disambiguate them, or…'. It follows from *Φ*(*s*) that *s* does not have any particular question for its content.

In contrast to the norm saying that one shouldn't *inquire* into pseudo-questions, (NAPS) doesn't just amount to the claim that one shouldn't do the impossible. For whereas it is impossible for one to actually inquire into a pseudo-question, it is not impossible for one to *try* to find the true answer to the question expressed by an *s* of the relevant sort *Φ*(*s*). Of course, where *Φ*(*s*), it follows that *there isn't* a question that is expressed by *s*. So it is impossible for one to actually find the true answer to the question expressed by *s*. But it is not impossible for one to *try* or to *attempt* to do that. In general, it is not impossible for one to try to do the impossible, even though it is of course impossible for one to do the impossible.

Now contrast (NAPS) to (IN). Whereas (IN) is used to assess attempts to *say* something meaningful, *ergo* attempts that need not have anything to do with inquiry (not necessarily), (NAPS) is used to assess attempts to find the true answers to purported questions, which are just attempts at successfully inquiring into questions. So (NAPS) is targeted at a kind of attempt that is much more intimately connected to inquiry, when compared to the kind of attempt that (IN) is targeted at.

To occupy oneself with a pseudo-question here is to attempt to find the true answer to the question expressed by a certain interrogative, to attempt to settle that purported question—but where there isn't such a question for one to settle or to find the

true answer to. For the target interrogative does not express any such particular question in the first place. An inquirer who abides by (NAPS) is an inquirer who does not make that kind of mistake.

§10.4 That the newly formulated norm is a norm of inquiry, though it is not a norm of inquiry into the properties and relations of linguistic expressions

The newly formulated norm (NAPS) sounds like a true norm. To the extent that it is concerned with the kind of attempt that is constitutive of the activity of inquiry—namely, that of trying to find the true answer to a question, or to settle a question—we can also call it a norm of inquiry.

(NAPS) is also an *instrumental* norm. The reason why one shouldn't search for the true answer to the question expressed by *s*, where *s* is a pseudo-question, is that such an attempt would be ineffective and counterproductive. For, again, the true answer to the question expressed by *s* does not exist, because *the* question expressed by *s* does not exist. So we justify (NAPS) in roughly the same way we have justified some of our prohibitive norms, such as the No False Presupposition Norm (NFP) and the Anti-Impossibility Norm (AIN). Since (NAPS) fills both of the conditions that are satisfied by the norms of our system—being concerned with inquiry and being an instrumental norm—we have then good reasons to include it in that system.

Just like we justify (NAPS) in roughly the same way we have justified some of our prohibitive norms, so we have to defend (NAPS) against the same kinds of objections that were raised against those norms.

More specifically, what if occupying oneself with a pseudo-question (in the present sense) is the means through which the inquirer comes to learn that *it is a pseudo-question* or that *it fails to express a question*? That surely looks like a benefit of her occupying herself with the pseudo-question, in which case it seems wrong to diagnose her activity as counterproductive.

But the response is that her goal-directed activity *is* counterproductive relative to the very aim or goal that is constitutive of it. The inquirer wanted to find the true answer to the question expressed by the target interrogative *s*. And she didn't find *that*—she couldn't possibly find it—for there is no true answer to question expressed by *s*. In fact, no activity is a means to that end. Any activity that the inquirer performs with the aim of achieving that end is misguided. That remains true even if, as contingency would have it, she benefits from her ineffective attempt in one way or another, which includes her learning that *s fails to express a question*, or that *s is a pseudo-question*.

The 'should' of (NAPS), as much as the 'should' of (AIN), (NFP), (INI), etc., is not an all-things-considered should, either. The fact that *there is no true answer to the question expressed by s* deprives the activity of attempting to find the true answer to

the question expressed by *s* from instrumental value, relative to the very goal that is constitutive of it. That activity is bound not to bring about its own end. This is compatible with granting that that activity has some other value for the agent.

The situation would be different if the inquirer were to inquire into the legitimate question of *whether s expresses a question*, say. There is nothing wrong, as far as (NAPS) goes, with the agent's inquiring into *whether s expresses a question*—which is a specific polar question, which admits of a unique true maximal answer. In order to issue a reproach by relying on (NAPS), however, we have to consider situations where the inquirer is simply searching for the true answer to the question expressed by an interrogative *s* such that *Φ*(*s*). In doing that, she is rather *assuming* or *presupposing* that there is a true answer to the question expressed by *s* (or at least her actions commit her to that claim), and therefore that there is a particular question expressed by *s*. She is then occupying herself with a pseudo-question. In contrast, when the inquirer is inquiring into the question of *whether s is a pseudo-question*, she is not occupying herself with a pseudo-question. She is rather occupying herself with a real question about interrogatives and their relations to questions.

When we assume that the inquirer is occupying herself with a pseudo-question, we are assuming that she is searching for the true answer to a question that is purportedly expressed by a certain interrogative, not that she is engaging in metalinguistic inquiry. To occupy oneself with a pseudo-question is not to perform legitimate inquiry into (real, not pseudo-) questions about the meanings of interrogatives. The norm that one shouldn't occupy oneself with pseudo-questions doesn't apply to the latter situation—it issues no reproach to the activity of inquiry that takes place in it. Neither does our explication of that norm, namely (NAPS), for (NAPS) is not a norm of inquiry into the properties and relations of linguistic expressions.

If, in trying to find the true answer to the question expressed by *s*, where *Φ*(*s*), the inquirer ends up discovering that *s fails to express a question*, then that is indeed a good consequence of her activity. But that good consequence would only have counted as an *achievement* of her goal-directed activity if what he had been trying to do was to settle the question of *whether s expresses a question*. By assumption, however, that is *not* what she was trying to do. A token of an interrogative *s* was produced, either by the inquirer herself or by someone else, and that prompted the inquirer to try to find the true answer to the question expressed by that token *s*. Since *Φ*(*s*), however, it follows that *s* conveys no question in particular. Whatever else the inquirer has found along the way, she did not find what she was searching for. She thinks or presumes that *s does* express a particular question. Since it doesn't, her search is just yet another wild-goose chase, beneficial to her as it may be in other respects.

§10.5 That those who occupy themselves with pseudo-questions may at the same time inquire into real questions

We have construed of the situation where the inquirer is occupying herself with a pseudo-question as follows: she is attempting to find the true answer to the question expressed by an interrogative *s,* but *s* actually fails to express a particular question, that is, *s* is such that *Φ*(*s*). And we also saw that, in attempting to find the true answer to the question expressed by *s*, the agent is not thereby inquiring into metalinguistic questions such as the question of *whether s expresses a question*, but is rather presupposing or acting as if *s* expresses one particular question.

In attempting to find the true answer to the question expressed by an interrogative *s* such that *Φ*(*s*), however, isn't the inquirer inquiring into some specific question, even though it is not one of those metalinguistic questions?

Let us look again at the situation where the inquirer is occupying herself with a pseudo-question in the present sense. The tokening of an interrogative *s* has prompted her to search for a true answer to some question, purportedly expressed by that occurrence of *s*. Give that *Φ*(*s*), it follows that *s* can either be disambiguated in different ways, or its meaning can be determined in different ways depending on how certain contextual parameters are fixed, etc.

But maybe the inquirer has thereby started inquiring into *one* of the particular questions that could have been expressed by that occurrence of *s*. Granted, the target token of *s* has failed to express any such question—but maybe the inquirer assumed that it did, and she went on to investigate into the question she assumed to have been expressed by it. If that is the case, then the inquirer who is trying to find the true answer to the question expressed by the target token of *s* is also legitimately inquiring into some real question—a candidate meaning of that token of *s*. She is occupying herself with a pseudo-question *and* at the same time occupying herself with a real question. Does that spell trouble for (NAPS)?

Let us pick an example to illustrate the phenomenon. We have given three examples of the raising of pseudo-questions above: one involving the utterance of 'Is he a fool?' at a context that fails to pick a referent for 'he', one involving the utterance of 'What is the meaning of life?' against a conversational background that offers no disambiguation of the expression 'meaning', and one involving the utterance of 'Does misleading higher-order evidence destroy epistemic justification?' against a conversational background that fails to specify how the technical term 'epistemic justification' is to be interpreted. Let us use the first one of these examples, on account of it being simpler than the other ones, but keeping in mind that similar conclusions would also apply to them.

So suppose again that a friend asks you ‘Is he a fool?’ while pointing in the direction of Daniels and O’Leary, but without doing enough (or even intending) to determine whether her use of ‘he’ refers to Daniels or rather to O’Leary. That token of ‘Is he a fool?’ expresses neither the question of *whether Daniels is a fool* nor the question of *whether O’Leary is a fool*. Yet, you start searching for the true answer to the question expressed by that token—and maybe you even tell your friend ‘I will tell you what the answer to that question is in a minute, let me think’. You then start thinking about *whether Daniels is a fool*. You somehow presume that your friend was not talking about O’Leary, maybe because you take it that O’Leary is clearly not a fool. You then retrieve the information you have about Daniels’ past behavior, the things he said and did, etc., to try and determine *whether Daniels is a fool*.

In this case, you are attempting to settle or to find the true answer to the question expressed by a token of ‘Is he a fool?’. But, in doing that, you are also inquiring into *whether Daniels is a fool*, which is a real, legitimate question. Now (NAPS) says you shouldn’t be making the former attempt—it is a waste of your time and resources, seeing as no particular question is actually expressed by the target occurrence of ‘Is he a fool?’ (compare to first asking your friend who she meant to refer to, which is not a waste of your time and resources). Doesn’t that mean, however, that (NAPS) thereby says you shouldn’t be inquiring into *whether Daniels is a fool*? That would be a bad result, seeing as (NAPS) is not designed to reproach inquiry into legitimate questions, but rather engagement with pseudo-questions.

It is more or less obvious that (NAPS) does *not* say or entail that you shouldn’t inquire into *whether Daniels is a fool*. In general, from the fact that one shouldn’t be such that φ, and that in satisfying φ one also satisfies ψ, it does not follow that one shouldn’t be such that ψ (clear case, where ‘S’ is again the ‘should’-operator: that $S\neg\varphi$ does not entail that $S\neg(\varphi \vee \neg\varphi)$—the latter is even absurd). In our example, (NAPS) issues a reproach to your attempt to find the true answer to the question expressed by a token of ‘Is he a fool?’, without issuing a reproach to facts that follow from or are constitutive of your so attempting—including both the fact that you’re still alive and the fact that you are inquiring into *whether Daniels is a fool*.

We can also think like this: if there is no token of some interrogative *s* such that the inquirer attempts to find the true answer to the question expressed by that token, then (NAPS) doesn’t issue a reproach to that inquirer. And it is possible for one to inquire into a question *Q* that could have been expressed by a token of *s without* attempting to find the true answer to the question expressed by a token of *s*. So (NAPS) doesn’t issue a reproach to an inquirer in the latter situation.

(NAPS) only comes into play when the inquirer is set in inquiring motion by the production or occurrence of a particular interrogative, to which she fruitlessly tries to be responsive to.

Chapter 11

§11.1 Where we present our system of norms of inquiry in a logically hierarchical order

After the last supplementation to our system of norms from the previous chapter—namely, the inclusion of the norm that tells us not to occupy ourselves with pseudo-questions (NAPS)—we can now present the main constituents of that system in a logically hierarchical order. That means presenting its more basic norms first, which can be conceived as *axioms* (albeit we have justified them outside the system), and afterwards some of their consequences, which can be conceived as *theorems*.

The basic norms or axioms are the following, leaving guise-parameters as well as other important parameter (such as time- and world-parameters) implicit:

(AIN) One shouldn't inquire into *Q* if it is impossible for one to know *Q*.

(INI) One shouldn't inquire into *Q* if one already knows *Q*.

(NAPS) One shouldn't attempt to find the true answer to the question expressed by an interrogative *s* such that *Φ*(*s*) (where '*Φ*' is an open sentence that characterizes an interrogative as a pseudo-question).

(GG) If one wants to know Q, but one's coarse-grained information is not yet enough to settle Q, then one should see to it that one gathers new information that bears on Q.

(GF) If one wants to know *Q*, and one's coarse-grained information *i* is enough to settle *Q*, then one should see to it that one competently deduces some fine-grained information *x* such that *v*(*x*) = *a*, where *a* is the maximal complete answer to *Q* such that $i \subseteq a$.

And some of the theorems or consequences of those five axioms include the following:

(G) If one wants to know *Q* then one should see to it that one inquires into *Q*.

(N) One shouldn't collect more information in an attempt to settle *Q* if one already knows *Q*.

(NUA) One shouldn't inquire into *Q* if none of the complete answers to *Q* are true.

(NFP) One shouldn't inquire into *Q* if *Q* has a false presupposition.

(ADN) One shouldn't inquire into *Q* while knowing that *Q does not have a true maximal answer*.

(K-AIM) One shouldn't inquire into *Q* if one knows that *it is impossible for one to know Q.*

(K-INI) One shouldn't inquire into *Q* if one knows that *one already knows Q*.

Norm (G), which we can simply call the 'Go Norm of Inquiry', is a consequence of (GG) and (GF) together, and it has gone unnoticed up to this point. It is easy to see how it follows from (GG) and (GF). Suppose the inquirer wants to settle *Q.* Either that inquirer's coarse-grained information *i* is enough to settle *Q*, or it isn't. If it is, then (GF) tells that inquirer to see to it that she inquires into *Q*, for it tells that inquirer to see to it that she deduces a fine-grained answer to *Q*. If it isn't, then (GG) tells that inquirer to see to it that she inquires into *Q*, for it tells that inquirer to see to it that she gather new information that bears on *Q*. Either way, then, assuming that the inquirer want to settle *Q*, she is directed to see to it that she inquires into *Q*. So (G).

Norm (N), as we saw, follows from (INI)—it is but an *instance* of (INI).

The No Untrue Answers Norm (NUA) and the No False Presupposition Norm (NFP) follow from the Anti-Impossibility Norm (AIN)—at least when we interpret the notion of possibility in (AIN) in the way we did, namely, as some kind of *circumstantial* possibility (§8.3). So does the Anti-Dissonance Norm (ADN), which itself follows from (NFP) (as seen in §7.2).

The epistemically qualified norms (K-AIM) and (K-INI), we also saw, follow from (AIN) and (INI), respectively, seeing as knowledge is factive (§9.3).

These are some of the theorems or consequences of the axioms of our system, then. We can regiment their derivation using standard deontic logic with a possible worlds semantics, treating our 'should'-operator as a kind of deontic necessity and requiring its accessibility relation to be at least serial.[119]

We have been somewhat sloppy with world- and time-parameters, and some of the consequences of our five axioms can only be proved once we make such

[119] See McNamara and Van De Putte (2022, §2) for an introduction and overview.

parameters explicit. We illustrate it with the following theorem (where 'NPW' abbreviates 'Norm of Possible Want'):

(NPW) One should be such that: one wants to settle *Q* at time *t* only if it is possible for one to know *Q* relative to some future time *t'* > *t*.

Let the notion of possibility in (NPW) be the same as the notion of possibility in (AIN). Now let us briefly see how this norm can be derived from our axioms.

If at time *t* it is true that one sees to it that one inquires into *Q*, then there is some future time *t'* > *t* such that one indeed inquires into *Q* at *t'*. Now suppose that (a) the inquirer wants to settle *Q* at *t,* and (b) that, at no future time *t'* > *t*, it is possible in the circumstantial sense of (AIN) for that inquirer to know *Q*. We are in fact supposing here that the inquirer violates (NPW). If from these assumptions it follows that either the inquirer has violated (G), or the inquirer has violated (AIN), then these two norms entail the norm that the inquirer should be such that: she wants to settle *Q* at *t* only if it is possible for her to know *Q* relative to some future time *t'* > *t*. That is, they entail (NPW).

Suppose, then, that our inquirer abides by (G) at *t*. Given assumption (a), namely, that she wants to know *Q* at time *t*, it follows that, at *t,* it is true that she sees to it that she inquires into *Q*. So for some future time *t'* > *t*, she does indeed inquire into *Q* at *t'*. But, by assumption (b), it is impossible for that inquirer to know *Q* at *t'*. So if our inquirer abides by (G) at *t*, then the fact that she violates (NPW) guarantees that she violates (AIN) at some such future time *t'*. If both (G) and (AIN) are true, then, it follows that (NPW) is true. Put another way, the inquirer can only abide by both (G) and (AIN) if she abides by (NPW).

(NPW) is but an instance of the general norm that one shouldn't want the impossible. There are many other consequences that can be derived from the axioms of our system—we have merely scratched the surface here. And similarly for the further consequences of the framework that we have adopted to theorize about inquiry here (which deploys formal constructs of questions and information in terms of sets of possible worlds and guises or representational vehicles). A more comprehensive exploration of all such consequences is left for future extensions of this work.

§11.2 Where we briefly state how we got here and conclude this investigation

How did we get to that system? We have started off this investigation by making certain assumptions about the nature of inquiry: that it is a kind of activity, and a goal-oriented one at that, that inquiry is always inquiry into some question or other, that its goal is to

settle some question or other, that certain questions can only be settled through the income of new information, whereas others can be settled from the armchair.

Then we refined those assumptions further, explicated the notions of coarse- and fine-grained information, offered a solution to a puzzle about armchair inquiry, and presented a model of inquiry that embodies that solution. We have thereby construed information-gathering inquiry as an activity designed to eliminate previous uneliminated possibilities, and armchair inquiry as an activity that is designed to fill gaps in the inquirer's representation of reality, or her picture of how the world is like. That led us to formulate and justify our pair of proactive norms, the Go Gather Norm (GG) and the Go Figure Norm (GG).

We have also argued that the constitutive goal of inquiry into a question *Q*, previously described as the goal of settling a question *Q*, is the goal of knowing *Q*. We argued for such a view in §§7.4–7.5. And that gave us an important premise with which to argue for the Ignorance Norm of Inquiry (INI), the Anti-Impossibility Norm (AIN), and also the Norm Against Our Occupying Ourselves With Pseudo-Questions (NAPS), all understood as instrumental norms of inquiry. (Importantly, we didn't accept the norm that *one shouldn't inquire into pseudo-questions*, instead of (NAPS), exactly because of our assumption that inquiry is always inquiry into some question or other).

The truth of our norms deploying the instrumental 'should' is grounded on three kinds of facts: (a) facts concerning the (circumstantial) possibility or impossibility of achieving the constitutive goal of inquiry, (b) or facts about whether it has already been achieved, (c) facts concerning which actions constitute the means to that goal. The norms from the system laid out above stem from the combination of the previously mentioned deliverances of our investigation with this very statement about what makes the norms of inquiry with the instrumental 'should' true.

We noted above that the exploration of further consequences of our system of norms and our theoretical framework is left for future work. But we din't mean to thereby imply that our very set of axioms cannot be further incremented. There should after all be other instrumental norms of inquiry that are *not* derivable from our axioms from above—though they may still be grounded on the assumptions and constructs of our theoretical framework. In particular, there should also be instrumental norms regarding which further questions one should inquire into, given the questions one is already inquiring into.[120] As far as what we have accomplished so far, the system of instrumental norms of inquiry is not fully fleshed out yet.

---

[120] Relatedly, see Whitcomb and Milson (forthcoming) for what they call 'norms of expansion for inquiring attitudes', and Rosa (2019, 2021) for norms about which questions one should suspend judgment about, given the questions one already suspends judgment about.

Whenever an inquirer fails to abide by any of our norms, we have grounds to criticize and correct them. Regarding norms that are sometimes hard to follow or to abide by, however, there is often a concern about how those norms can be *guiding*, that is, of how an agent can actually deploy them in thought and action. And it might be thought that *it is* hard for an inquirer to abide by some of our norms. So how can an inquirer by guided by those norms?

Well, the inquirer has to ask herself: Does it look like this is the kind of question that can be settled from the armchair? Or is this the kind of question that requires more data to be settled? Does it look like I know the answer to this question already? Does it look like it is *possible* for me to know its true answer? What are the typical signs of unknowability? And does it look like a pseudo-question has been raised? What are the signs of *that*? She will then be able to use her available evidence concerning such questions, coupled with our norms, to decide what to do—to inquire or not to inquire, and in what way.

References

Archer, Avery (2018). ‘Wondering about What You Know’, *Analysis* 78(4): 596–604.

Bach, Kent (2001). ‘Speaking loosely: Sentence nonliterality’, *Midwest Studies in Philosophy* 25(1): 249–263.

Baker, Derek (2018). ‘Skepticism about Ought *Simpliciter*’, in R. Schafer-Landau (ed.) *Oxford Studies in Metaethics,* Volume 13, Oxford: Oxford University Press, pp. 230–252.

Bar-Hillel, Yehoshua and Rudolph Carnap (1952). *An Outline of a Theory of Semantic Information*, rep. in Bar-Hillel (1964) (ed.), Language and Information, Addison-Wesley, Reading, MA, pp. 221Y274.

Bechtel, William (2016). ‘Investigating neural representations: the case of place cells’, *Synthese* 193(5): 1287–1321.

Belnap, Nuel (1983). ‘Approaches to the Semantics of Questions in Natural Language (I)’, in Bäuerle, R., Schwarze, C., von Stechow (Des.) *Meaning, Use, and Interpretation of Language*, De Gruyter, pp. 22–29.

Belnap, Nuel and Thomas Steel (1976). *The Logic of Questions and Answers*, New Haven CT: Yale University Press.

Bernecker, Sven (2014). ‘How to understand the extended mind’, *Philosophical Issues* 24(1): 1–23.

Blome-Tillmann, Michael (2014). *Knowledge and Presuppositions*, New York: Oxford University Press.

Boghossian, Paul and Timothy Williamson (2020). *Debating the A Priori*, Oxford: Oxford University Press.

Borgoni, Cristina, Dirk Kindermann and Andrea Onofri (2021) *The Fragmented Mind*, Oxford: Oxford University Press.

Braine, M. D. S. & O’Brien, D. P. (1998). *Mental Logic*, Mahwah NJ: Erlbaum.

Brunero, John (2010). ‘The scope of rational requirements’, *The Philosophical Quarterly* 60(238): 28–49.

Camp, Elizabeth (2007). ‘Thinking with maps’, *Philosophical Perspectives* 21(1): 145–182.

Carnap, Rudolph (1947). *Meaning and Necessity: A Study in Semantics and Modal Logic*, Chicago IL: University of Chicago Press.

Carnap, Rudolf and Yehoshua Bar-Hillel (1952). ‘An Outline of a Theory of Semantic Information’, Technical report 247, Cambridge, MA: Research Laboratory of Electronics, Massachusetts Institute of Technology.

Carruthers, Peter (2018). ‘Basic questions’, *Mind and Language* 33(2): 130–147.

Carter, Sam (2021). ‘The dynamics of loose talk’, *Noûs* 55(1): 171–198.

Casullo, Albert and J. C. Thurow, (2013). *The A Priori in Philosophy*, Oxford: Oxford University Press.
Ciardelli, Ivano, Jeroen Groenendijk and Floris Roelofsen (2019). *Inquisitive semantics*, Oxford: Oxford University Press.
Ciardelli, Ivano (2021). 'Why we need a question semantics', in M. Cordes (ed.) *Asking and Answering*, Tübingen: Narr Francke Attempto, pp. 15–47.
Clark, Andy and David Chalmers (1998). 'The extended mind', *Analysis* 58(1): 7–19.
Cohen, Stewart (1988). 'How to be a fallibilist', *Philosophical Perspectives* 2: 91–123.
Coppin, Peter, Abrose Li and Michael Carnevale (2016). 'Iconic Properties are Lost when Translating Visual Graphics to Text for Accessibility'. URL: <http://openresearch.ocadu.ca/id/eprint/1035/>.
Cresswell, M. J. (1975). 'Hyperintensional Logic', *Studia Logica* 34(1): 25–38.
Cresswell, M. J. (1985). *Structured Meanings*, Cambridge, MA: MIT Press.
Cummins, Robert (1989). *Meaning and Mental Representation*, Cambrige, MA: MIT Press.
Davidson, Donald (1968). 'On saying that', *Synthese* 19(1/2): 130–146.
Davis, Martin ([1958] 1982). *Computability & Unsolvability*, New York: Dover Publications.
De Morgan, Augustus (1847). *Formal Logic*, London: Taylor & Walton.
Dennett, Daniel (1987). *The Intentional Stance*, Cambridge MA.: The MIT Press.
Dennett, Daniel (2002). 'True Believers: The Intentional Strategy and Why It Works', in D. Chalmers (ed.) *Philosophy of Mind—Classic and Contemporary Readings*, New York: Oxford University Press, pp. 556–568.
Dewey, John (1938). *Logic: The Theory of Inquiry*, New York: Henry Holt and Company.
Dogramaci, Sinan (2018). 'Rational credence through reasoning', *Philosophers' Imprint* 18(11): 1–25.
Dowell, Janice (2017). 'Contextualism about epistemic modals', in J. Jenkins Ichikawa (ed.) *The Routledge Handbook of Epistemic Contextualism*, New York: Routledge, pp. 388–399.
Dror, Shany, Andrea Sommese, Ádám Miklósi, Andrea Temesi and Claudia Fugazza (2022). 'Multisensory mental representation of objects in typical and Gifted Word Learner dogs', *Animal Cognition* 25: 1557–1566.
Dummett, Michael (1978). 'The justification of deduction', in M. Dummett, *Truth and other enigmas*, Cambridge, MA: Harvard University Press, pp. 166-185.
Eder, Anna-Maria (2019). 'Evidential Probabilities and Credences', *British Journal for the Philosophy of Science* 74(1): 1–21.

Egan, Frances (2020). ‘A Deflationary Account of Mental Representation’, in J. Smortchkova, K. Dolega and T. Schlicht (eds.) *Mental Representations*, New York: Oxford University Press, pp. 26–53.
Elga, Adam and Agustín Rayo (2021). ‘Fragmentation and Information Access’, in C. Borgoni, D. Kindermann and A. Onofri (eds.) *The Fragmented Mind*, Oxford: Oxford University Press, pp. 37–53.
Falbo, Arianna (2023). ‘Inquiring Minds Want to Improve’, *Australasian Journal of Philosophy* 101(2): 298–312.
Fantl, Jeremy and Matthew McGrath (2002). ‘Evidence, Pragmatics and Justification’, *The Philosophical Review* 111(1): 67–94.
Feldman, Richard (1985). ‘Reliability and Justification’, *The Monist* 68(2): 159–174.
Feldman, Richard and Earl Conee (1985). ‘Evidentialism’, *Philosophical Studies* 48: 15–34.
Fine, Kit (1975). ‘Vagueness, Truth and Logic’, *Synthese* 30: 265–300.
Foot, Philippa (1972). ‘Morality as a System of Hypothetical Imperatives’, *The Philosophical Review* 81(3): 305–316.
Friedman, Jane (2013). ‘Question-Directed Attitudes’, *Philosophical Perspectives* 27: 145–174.
Friedman, Jane (2017). ‘Why suspend judging?’, *Noûs* 51(2): 302–326.
Friedman, Jane (2019). ‘Inquiry and Belief’, *Noûs* 52(2): 249–510.
Friedman, Jane (2020). ‘The epistemic and the zetetic’, *Philosophical Review* 129(4): 501–536.
Friedman, Jane (forthcoming). ‘Zetetic Epistemology’, in B. Reed and A. Flowerree (eds.) *Towards an Expansive Epistemology: Norms, Action, and the Social Sphere*, Routledge.
Goel, Vinod (2005). ‘Cognitive Neuroscience of Deductive Reasoning’, in K.J. Holyoak and R.G. Morrison (eds.) *The Cambridge Handbook of Thinking and Reasoning*, New York: Cambridge University Press, pp. 475-492.
Greco, Daniel (2014). ‘Could KK be OK?’, *The Journal of Philosophy* 111(4): 169–197.
Greco, Daniel (forthcoming). ‘Justifications and Excuses in Epistemology’, *Noûs*, DOI: 10.1111/nous.12309.
Greenspan, Patricia (1975). ‘Conditional Oughts and Hypothetical Imperatives’, *The Journal of Philosophy* 72(10): 259–276.
Groenendijk, Jeroen and Martin Stokhof (1982), ‘Semantic analysis of wh-complements’. Linguistics and Philosophy 5:175–233.
Groenendijk, Jeroen and Stokhof, Martin (1984). *Studies on the semantics of questions and the pragmatics of answers*. Dissertation, University of Amsterdam, Amsterdam.

Grundmann, Thomas (2022). ‘Dependent Reliability: Why And How Conditional Reliability Should Be Replaced By It’, *Philosophy and Phenomenological Research* 105(1): 144–159.
Hamblin, C. L. (1973). ‘Questions in Montague English’, *Foundations of Language* 10(1): 41– 53.
Haziza, Eliran (2023). ‘Norms of Inquiry’, *Philosophy Compass* 18(12): e12952.
Hawthorne, John (2004). *Knowledge and Lotteries*, New York: Oxford University Press.
Hintikka, Jaakko (1962). *Knowledge and Belief—An Introduction to the Logic of the Two Notions*, Cornell University Press.
Hintikka, Jaakko (1973). *Logic, Language-Games and Information*, Clarendon: Oxford University Press.
Howard-Snyder, Daniel (2013). ‘Propositional faith: what it is and what it is not’, *American Philosophical Quarterly* 50(4): 357–372.
Hutto, Daniel and Eric Myin (2014). ‘Neural representations not needed—no more pleas, please’, *Phenomenology and the Cognitive Sciences* 13(2): 241–256.
Jago, Mark (2013). ‘The Content of Deduction’, *Journal of Philosophical Logic* 42: 317-334.
Johnson-Laird, Philip and Ruth Byrne (1995). ‘A model point of view’, *Thinking and Reasoning* 1: 339–350.
Johnson-Laird, Philip and Sangeet Khemlani (2013). ‘Toward a unified theory of reasoning’, *Psychology of Learning and Motivation* 59: 1-42.
Kant, Immanuel ([1781, 1787] 1998). *Kritik der reinen Vernunft*, J. Timmermann (ed.), Hamburg: Felix Meiner Verlag.
Kaplan, David (1989). ‘Demonstratives’, in J. Almog, J. Perry and H. Wettstein (eds.), *Themes from Kaplan*, New York: Oxford University Press, pp. 481–563.
Karttunen, Lauri (1977). ‘Syntax and semantics of questions’, *Linguistics and Philosophy* 1: 3–44.
Keefe, Rosanna (2003). *Theories of Vagueness*, Cambridge: Cambridge University Press.
Kelp, Christoph (2018). ‘Inquiry, knowledge and understanding’, Synthese, 198(7): 1583–1593.
Kelp, Christoph (2021). ‘Theory of inquiry’, *Philosophy and Phenomenological Research* 103(2): 359–384.
Kleene, S. C. (1952). *Introduction to Metamathematics*, Amsterdam: North Holland.
Kolodny, Niko (2005). ‘Why Be Rational?’, *Mind* 114(455): 509–563.
Kolodny, Niko (2007). ‘How Does Coherence Matter?’, *Proceedings of the Aristotelian Society* 107: 229–263.

Kolodny, Niko (2018). ‘Instrumental reasons’, in D. Star (ed.) *The Oxford Handbook of Reasons and Normativity*, Oxford: Oxford University Press.
Kratzer, Angelika (1977). ‘What ‘Must’ and ‘Can’ Must and Can Mean’, *Linguistics and Philosophy* 1: 337–355.
Kratzer, Angelika (1981). ‘The Notional Category of Modality’, in H. J. Eikmeyer and H. Rieser (eds.), *Words, Worlds, and Contexts*, Berlin: W. de Gruyter, pp. 38–74.
Kratzer, Angelika (2022). ‘Attitude Ascriptions and Speech Reports’, in D. Altshuler (ed.) *Linguistics Meets Philosophy*, Cambridge: Cambridge University Press, pp. 17–49.
Kvanvig, Jonathan (1998). ‘Why Should Inquiring Minds Want to Know?’, *The Monist* 81(3): 426–451.
Landy, David, Colin Allen and Carlos Zednik (2014). ‘A perceptual account of symbolic reasoning’, *Frontiers in Psychology* 275(5): 1–10.
Larkin, Jill and Herbert Simon (1987). ‘Why a Diagram is (Sometimes) Worth Ten Thousand Words’, *Cognitive Science* 11(1): 65–100.
Lasersohn, Peter (1999). ‘Pragmatic halos’, *Language* 75(3): 522–551.
Lepore, Ernest and Barry Loewer (1989). ‘You can say that again’, *Midwest Studies in Philosophy* 14: 338–356.
Lewis, David (1973). *Counterfactuals*, Cambridge MA: Harvard University Press.
Lewis, David (1974). ‘Radical Interpretation’, *Synthese* 27: 331–344.
Lewis, David (1996). ‘Elusive Knowledge’, *Australasian Journal of Philosophy* 74(4): 549–567.
Łukasiewicz, Jan (1967). ‘Philosophical Remarks on Many-Valued Systems of Propositional Logic’, in S. McCall (ed.) *Polish Logic*, 1920–1939, Oxford: Clarendon Press, pp. 228–251.
Luper, Steven (2020). ‘Epistemic Closure’, in E. N. Zalta (ed.) *The Stanford Encyclopedia of Philosophy*, Summer 2020 edition, URL: <https://plato.stanford.edu/archives/sum2020/entries/closure-epistemic/>.
Manktelow, Ken (1999). *Reasoning and Thinking*, East Sussex: Taylor & Francis.
Martínez, Manolo and Marc Artiga (forthcoming). ‘Neural Oscillations as Representations’, *The British Journal for the Philosophy of Science*, DOI: 10.1086/714914.
McCain, Kevin (2014). *Evidentialism and Epistemic Justification*, New York: Routledge.
McNamara, Paul and Frederik Van De Putte (2022). ‘Deontic Logic’, in E. N. Zalta (ed.) *The Stanford Encyclopedia of Philosophy*, Fall 2021 edition, URL: <https://plato.stanford.edu/archives/fall2022/entries/logic-deontic/>.
Menzel, Christopher (2021). ‘Possible Worlds’, in E. N. Zalta (ed.) *The Stanford Encyclopedia of Philosophy*, Fall 2021 edition, URL: <https://plato.stanford.edu/archives/fall2021/entries/possible-worlds/>.

Moss, Sarah (2018). *Probabilistic Knowledge*, New York: Oxford University Press.
Moss, Sarah (2019). 'Full Belief and Loose Speech', *The Philosophical Review* 128(3): 255–291.
Pacuit, Eric (2013a). 'Dynamic Epistemic Logic I: Modelling Knowledge and Belief', *Philosophy Compass* 8(9): 798–814.
Pacuit, Eric (2013b). 'Dynamic Epistemic Logic II: Logics of Information Change', *Philosophy Compass* 8(9): 815–833.
Pavese, Carlotta and Beddor, Bob (forthcoming). 'Skills as Knowledge', *Australasian Journal of Philosophy*, DOI: 10.1080/00048402.2022.2056753.
Portner, Paul (2009). *Modality*, New York: Oxford University Press.
Quine, W. V. O. (1960). *Word and Object*, Cambridge MA: MIT Press.
Ramsey, William (2007). *Representation Reconsidered*, Cambridge: Cambridge University Press.
Rips, Lance J. (1994). *The psychology of proof*, Cambridge MA: The MIT Press.
Roelofsen, Floris (2019). 'Semantic theories of questions', in M. Aronoff (ed.) *Oxford Encyclopedia of Linguistics*, Oxford University Press. URL: https://doi.org/10.1093/acrefore/9780199384655.013.504.
Rosa, Luis (2019). 'Logical Principles of Agnosticism', *Erkenntnis* 84(6): 1263–1283.
Rosa, Luis (2021). 'Rational requirements for suspended judgment', *Philosophical Studies* 178(2): 385–406.
Rosa, Luis (2022). 'Coherence and Knowability', *The Philosophical Quarterly* 72(4): 960–978.
Rosa, Luis (forthcoming). 'Doubt', in K. Sylvan (ed.) *The Blackwell Companion to Epistemology*, *3rd Edition*, Blackwell Publishing.
Saha, Avishkar, O. M. Maldonado, Chris Russell and Richard Bowden (2021). 'Translating Images into Maps', URL: <https://doi.org/10.48550/arXiv.2110.00966>.
Salmon, Nathan (1995). 'Being of Two Minds: Belief with Doubt', *Noûs* 29(1): 1–20.
Schaffer, Jonathan (2005). 'Contrastive Knowledge', in T. S. Gendler & J. Hawthorne (eds.), *Oxford Studies in Epistemology 1*, Oxford University Press, pp. 235–270.
Schroeder, Mark (2012). 'Stakes, withholding, and pragmatic encroachment on knowledge', *Philosophical Studies* 160(2): 265–285.
Sequoiah-Grayson, Sebastian (2007). 'The Scandal of Deduction', *Journal of Philosophical Logic* 37(1): 67–94.
Shea, Nicolas (2018). *Representation in Cognitive Science*, Oxford: Oxford University Press.
Staffel, Julia (2013). 'Can there be reasoning with degrees of belief?', *Synthese* 190(16): 3535–3551.

Staffel, Julia (2019). ‘Credences and suspended judgments as transitional attitudes’, *Philosophical Issues* 29(1): 281–294.
Stalnaker, Robert (1970). ‘Pragmatics’, *Synthese* 22(1/2): 272–289.
Stalnaker, Robert (1984). *Inquiry*, Cambridge MA: MIT Press.
Stalnaker, Robert (1991). ‘The Problem of Logical Omniscience, I’, *Synthese* 89(3): 425–440.
Stanley, Jason (2008). ‘Knowledge and certainty’, *Philosophical Issues* 18(1): 35–57.
Stenning, Keith and van Michiel Lambalgen (2008). *Human Reasoning and Cognitive Science*, Cambridge, MA: MIT Press.
Tarski, Alfred ([1953] 2010). *Undecidable Theories: Studies in Logic and the Foundation of Mathematics*, New York: Dover Publications.
Thomasson, Amie L. (2009). ‘Answerable and Unanswerable Questions’, in D. Chalmers, D. Manley and R. Wasserman (eds.) *Metametaphysics: New Essays on the Foundations of Ontology*, Oxford: Oxford University Press, pp. 444–471.
Thomson, Eric and Gualtiero Piccinini (2018). ‘Neural Representations Observed’, *Minds and Machines* 28(1): 191–235.
Thorstad, David (2021). ‘Inquiry and the epistemic’, *Philosophical Studies* 178(9): 2913–2928.
Tiffany, Evan (2007). ‘Deflationary Normative Pluralism’, *Canadian Journal of Philosophy* 37(5): 231–262.
Tye, Michael (1994). ‘Sorites paradoxes and the semantics of vagueness’, *Philosophical Perspectives* 8: 189–206.
van Benthem, Johan and Fernando R. Velázquez-Quesada (2010). ‘The dynamics of awareness’, *Synthese* 177(1): 245–267.
van Ditmarsch, Hans, Wiebe van der Hoek and Barteld Kooi (2008). *Dynamic Epistemic Logic* (Synthese Library), Dordrecht: Springer.
van Elswyk, Peter and Yasha Sapir (2021). ‘Hedging and the ignorance norm on inquiry’, *Synthese* 199(3): 5837–5859.
van Fraassen, Bas (1980). *The Scientific Image*, Oxford: Oxford University Press.
Way, Jonathan (2010). ‘Defending the wide-scope approach to instrumental reason’, *Philosophical Studies* 147(2): 213–233.
Way, Jonathan (2011). ‘The symmetry of rational requirements’, *Philosophical Studies* 155(2): 227–239.
Whitcomb, Dennis (2010). ‘Curiosity was framed’, *Philosophical Quarterly* 81(3): 664–687.
Whitcomb, Dennis (2017). ‘One kind of asking’, *Philosophical Quarterly* 67(266): 148–168.

Whitcomb, Dennis and Jared Millson (forthcoming). ‘Inquiring Attitudes and Erotetic Logic: Norms of Restriction and Expansion’, *Journal of the American Philosophical Association*. DOI: 10.1017/apa.2023.13.
Willard-Kyle, Christopher (forthcoming). ‘The Knowledge Norm for Inquiry’, *Journal of Philosophy*.
Williamson, Timothy (2000). *Knowledge and its Limits*. Oxford: Oxford University Press.
Williamson, Timothy (2020). *Suppose and Tell*. Oxford: Oxford University Press.
Williamson, Timothy (forthcoming). ‘Justification, Excuses, and Sceptical Scenarios’, In F. Dorsch and J. Dutant (eds.) *The New Evil Demon*, Oxford University Press.
Wittgenstein, Ludwig (1922). *Tractatus Logico-Philosophicus*, London: Kegan Paul.
Woodard, Elise (forthcoming). ‘Why Double-Check?’ *Episteme*.
Worsnip, Alex (2021). *Fitting Things Together: Coherence and the Demands of Structural Rationality*, New York: Oxford University Press.
Yap, Audrey (2014). ‘Idealization, epistemic logic and epistemology’, *Synthese* 191(14), 3351–3366