Chapter 4:

Reformulation as a Constitutive Aim of Science

4.1 Introduction

In developing constructive empiricism, van Fraassen (1980) framed the debate over scientific realism as a disagreement about the aims of science. Scientific realism takes the aim of science to be approximate truth, including claims about physical unobservables.[1] In contrast, constructive empiricism defends a weaker aim for science: empirical adequacy. Empirical adequacy requires truth only about observable states of affairs. Nevertheless, as Rosen (1994) pointed out, the foundational status of these "aims of science" is not clear. Are they descriptive claims about scientific sociology? Prescriptive claims about what scientists ought to aim at, independently of their current goals? Or are they merely a useful way of talking, perhaps a kind of fictional story about science?

I will analyze the aims of science by focusing on those aims that are constitutive of scientific activity. Constitutive aims define the *minimal* criteria of success, i.e. the criteria that must be met for scientific activity to succeed at all. Unlike success criteria in general, constitutive features are often hidden from participants. For instance, native language speakers are often unaware of the grammatical rules they follow. I will argue that one constitutive aim of science is the ability to solve all physically possible empirical problems. I will call this aim *problem-solving adequacy*. My proposal is similar to Laudan's, who argues that "the aim of science is to secure theories with a high problem-solving effectiveness" (1996, p. 78). Of course, realists and antirealists disagree about the minimal success criteria for solving a scientific problem. Whereas realists posit truth, construc-

[1]Philosophers commonly speak about the aims of science, especially physics. See, for instance, Loewer (2007, pp. 319, 322, 326), who in passing distinguishes the aims of physics from the aims of metaphysics, or Potochnik (2017). Earman and Roberts "presuppose that science does aim to discover laws (among lots of other things)" (2005, p. 254). Section 4.8 considers objections to speaking about aims of science.

tive empiricists posit empirical adequacy. Because my account enriches both realism and antirealism, I will remain neutral on this contentious debate here.[2]

To argue that problem-solving adequacy is a constitutive aim of science, I focus on a ubiquitous aspect of scientific practice: reformulating scientific theories and problem-solving procedures. Unamended, both realism and antirealism leave mysterious why scientists go to such great lengths to clarify which epistemic resources they need to solve problems, even when this clarification does not make their theory any more empirically adequate or true. I will argue that beyond seeking truth or empirical adequacy, scientists ought to clarify the epistemic structure of their theories. This constitutive aim arises from the need to prepare for any possible problem-solving context that could arise. My argument shows that reformulations are not merely instrumentally valuable for either truth or empirical adequacy. Instead, reformulations are valuable for their own sake, since they are constitutive of clarifying epistemic structure, which is itself a constitutive aim of science.

The argument of this chapter rebuts a worry raised in Chapter 1. There, I wondered whether the value of reformulating might ultimately just be instrumental. If so, then conceptualism would itself be a kind of instrumentalism about reformulations (Section 1.4). Here, I show that conceptualism does not reduce to instrumentalism. Reformulations possess a kind of final epistemic value that goes beyond their instrumental value.

4.2 Constitutive Aims

Empiricism and Scientific Aims

Constructive empiricism and scientific realism posit the same truth conditions for a scientific theory: a theory is true only if it matches reality at all length scales.[3] Furthermore, constructive empiricists agree with realists about the meaningfulness of claims about physical unobservables (even ones that are in-principle empirically inaccessible). For these reasons, constructive empiricism is not a variety of verificationism. Instead,

[2] Elsewhere, I argue that empirical adequacy, rather than truth, is the correct minimal success criterion for solving a physical problem. Since agents lack any competence for perceiving unobservables, the constitutive aim of perception is accuracy about observables. Developing this argument would take us too far afield here.

[3] Traditionally, these positions rely on a correspondence theory of truth.

constructive empiricism restricts the aim of science from truth at all length scales to truth about observables. Aiming for truth about observables is equivalent to aiming for empirical adequacy.

More precisely, a theory is empirically adequate provided that it accurately describes and predicts all true claims about observables.[4] Colloquially, an empirically adequate theory 'saves the phenomena.' Although less demanding than achieving truth, empirical adequacy remains a difficult aim to satisfy. A theory is empirically adequate only if it accurately represents the truth about observables past, present, and future, including not just what will be observed, but what could be observed. As I discuss in Section 4.4, empirical adequacy—like truth—is best understood as a *futuristic* aim: we will plausibly never obtain an empirically adequate theory. Rather, we approach greater empirical adequacy over time.

Rosen (1994) points out that as a sociological matter, many (if not most) scientists are aiming for more than empirical adequacy. Hence, for constructive empiricism to be plausible, it must distinguish the aims of science from the aims of individual scientists, even their aggregate aims. Indeed, van Fraassen contends that the aims of science are not the same as the aims of most or even all scientists (1994, p. 185). This generates a distinction between intentionality at the level of agents vs. intentionality at the level of the activity of science. Intentionality at the agent-level includes the intentions, opinions, and motivations of individual scientists. Intentionality at the activity-level abstracts away from these individual intentions. The aims of science exist at the activity-level.

To insulate the aims of science from the intentions of individual scientists, van Fraassen characterizes aims as criteria of success (1980, p. 8). As he notes, criteria of success often exist independently of the personal aims or motivations for engaging in a given activity. As a guiding analogy, both van Fraassen and Rosen consider the game of chess. One aim of chess is to checkmate your opponent. This aim serves as a shared criterion of success, independently of the particular reasons why anyone is playing chess. It is in this sense that van Fraassen claims that the "criterion of success in actual practice is empirical adequacy," regardless of scientists' individual aims, motives, or intentions

[4]On the semantic (model-theoretic) view of theories, this requires that the empirical substructure of the theory's models is isomorphic to observable reality (van Fraassen 1980, p. 64). However, I intend to remain neutral on the debate between syntactic vs. semantic approaches to scientific theories, especially given recent arguments that they are not competitors (Halvorson 2016; Lutz 2017).

(1994, p. 182). It is at least plausible that empirical adequacy could be an aim of science at the activity-level without being an explicit aim at the agent-level.

Nevertheless, van Fraassen's proposal is too weak to exclude truth from being an aim of science. If the aims of science are merely criteria of success, there is no reason why truth cannot be an aim as well. For surely, obtaining truth about unobservables would be a great success, in addition to empirical adequacy. Hence, interpreting the aims of science as criteria of success does not favor constructive empiricism over scientific realism. Similarly in chess, a player would demonstrate supreme mastery if they were able to not only checkmate their opponent but also capture all of their pieces along the way (or if they were able to checkmate their opponent as quickly as possible). Noting that checkmating your opponent is a criterion for success does not exclude other criteria for success. Likewise, recognizing that empirical adequacy is a criterion for success does not exclude either aiming at the truth or aiming at fundamental structure.

Van Fraassen's proposal faces a further, related challenge. Rosen wonders why the aims of science—defined as criteria of success—are not themselves "constituted by the conscious understandings of the participants" (1994, p. 146). Why isn't it the case that scientists' own understanding of science determines what qualifies as success? In response, van Fraassen grants that the conscious understandings of scientists define the activity 'science,' but that nevertheless this "does not mean what all the participants say they are doing is what they are doing" (1994, p. 186). This response grants the possibility that all scientists aim at truth, without truth being a criterion of success for science. But this response is implausible for success criteria in general. If the majority of scientists decided that the gold standard of experimental design requires insulating all measurement devices from Wi-Fi signals, this would generate a success criterion. It is therefore implausible that success criteria in general are independent from the conscious understandings of scientists.

Aims that are Constitutive

Fortunately for the aspiring empiricist, van Fraassen's response stands a chance when restricted to a more narrow class of success criteria. I will call these *constitutive aims* of an activity. They define *minimal* criteria of success. In chess, a constitutive aim is to checkmate your opponent: this is not merely a criterion for success; it is a minimal criterion.[5] In contrast, capturing all of your opponents' pieces—while no doubt grounds for success—is not constitutive of this success. Even if all chess players and chess federations agreed tomorrow that winners would receive twice the prize money for capturing all pieces, this would not make capturing all pieces into a constitutive aim. Capturing all pieces would still be unnecessary for winning a game of chess.

To understand what it takes for a criterion to be minimal, I invoke the notion of a constitutive feature. 'Constitutive aims' are then *the constitutive criteria for success*, determined by the activity itself. A *constitutive feature* is a feature something has in virtue of being itself. A thing's constitutive features are its conceptually essential or indispensable aspects. For instance, constitutive of the bishop piece in chess is moving along diagonals. Constitutive features, including constitutive rules, stand in contrast to arbitrary conventions, such as designating a particular piece as 'the king' (Searle 1995, p. 49). Instead, they are connected to the very possibility of the object or activity. As Max Black notes, "Chess played for the sake of losing is not chess but some other game" (1970, p. 159).[6]

In his accounts of perception, action, and knowledge, Sosa (2015) invokes a similar notion of a constitutive aim: a performance succeeds provided that it attains its constitutive aim. Moreover, attaining a constitutive aim does not preclude a performance from succeeding "even more fully" (2015, p. 14). It is in this sense that I take constitutive aims to define criteria for minimal success.[7]

Unlike aims in general, constitutive aims have the features that van Fraassen's account requires. Constitutive features, rules, or aims are often not manifest or perspicuous to participants. For instance, native speakers of English are often unaware of its grammatical rules, even rules that they routinely follow. Their lack of awareness is not evidence that there are no grammatical rules.[8] Moreover, English speakers could collectively decide to institute new rules, but this would not necessarily affect constitutive rules of their

[5]Similarly, we can understand a constitutive aim of an individual chess move as being "to do what will best help the player towards winning, or at least toward averting immediate defeat" (Sosa 2015, p. 126).

[6]We also commonly talk about constitutive norms. In her discussion of Peirce, Misak claims that "it is a constitutive norm of belief that a belief is responsive to the evidence and argument for or against it" (2013, p. 35).

[7]Sosa prefers to say that agents *succeed fully* when they attain a constitutive aim, but I worry this leads to grammatical confusion: in many contexts, it seems strange to say that "you can succeed fully even if you might have succeeded even more fully" (Sosa 2015, p. 14). The notion of minimal success preserves the relevant distinction while avoiding this infelicity.

[8]For a similar discussion in the context of semantic rules, see Thomasson (2020, p. 65).

language. As another example, many people are unaware that the value of money is a social construct, but this does not make money's value any less socially constructed. People might even believe money is valuable due to holding a false theory, such as that it must be "backed by gold" (Searle 1995, p. 47). The upshot is that scientists may very well be unaware of or even mistaken about the constitutive aims of science, without their lack of understanding affecting these aims. Focusing on constitutive aims thereby clarifies why aims-talk is not inherently sociological. Constitutive aims support a distinction between the intentional aspects of science vs. the intentions or opinions of individual scientists.

In the next section, I propose a method for identifying constitutive aims of science. Before that, two points of clarification are in order. First, the constitutive aims of an activity need not be sufficient for defining that activity. Multiple different games might all have a constitutive aim of checkmating opponents. For this reason, the constitutive aims that I defend for science do not demarcate science from either pseudoscience or non-science. If they are honest investigators, cryptozoologists searching for Bigfoot are bound by empirical adequacy. Likewise, ordinary problem-solving (such as figuring out how to cook pasta) involves considering possible scenarios and striving for empirical adequacy. But we wouldn't count cooking pasta as science, unless it were investigated sufficiently systematically![9]

Second, science is typically a community activity. Hence, the normative labor of constitutive aims is distributed: scientific agents ought to desire that members of their epistemic community are collectively working toward satisfying all of the constitutive aims of science. Each individual agent does not shoulder the normative weight of these obligations themselves. For whatever reasons, there may be particular tasks that individual agents abhor. They simply ought to desire that some magnanimous colleague will someday pick up the slack. The constitutive aims of science thereby generate *pro tanto* reasons for action. Such reasons can be outweighed at the individual level by other considerations, but never canceled.

[9] I thank Angela Sun and Sumeet Patwardhan for raising these demarcation issues.

4.3 Scientific Planning as Problem-Solving

Isolating constitutive aims of science requires identifying constitutive features of scientific activity. I will argue that one relevant constitutive feature is *scientific planning*, consisting of empirical problem-solving. Essential to scientific activity is answering questions about the physical world. As a method for answering empirical questions, science takes many forms: describing the past (e.g. archaeology), predicting the future (e.g. meteorology), building devices, synthesizing substances, etc. In different ways, each of these scientific activities contributes different methods for planning about how to interact with the world. They are methods for planning how to solve empirical problems, i.e. problems about the physical world. We can therefore understand one constitutive activity of science as being a form of planning. Scientists are planning how to solve problems, which amounts to planning how to answer their questions about the physical world. Even in the historical sciences, empirical problem-solving involves planning how to gather and assess evidence about the past.

Given this constitutive activity, we can determine its constitutive goals. Doing so will take us closer to isolating minimal success criteria for science. Beginning with local instances of scientific planning, we simply need to determine the local constitutive goals. Since this planning amounts to trying to solve an empirical problem, the constitutive goal is evidently to *solve the problem*. The foundational question is thereby to determine the minimal success criteria for solving an empirical problem: What makes a putative answer to a scientific question minimally successful? Answering this question helps identify local constitutive aims of science. We can then generalize from these local aims to determine global constitutive aims of science.

Following Goldman (1986, pp. 126–127), we can always define a problem in terms of asking a question, and we can define a solution in terms of an answer to that question. Schematically, an agent S has a scientific problem Q if and only if Q is a question about the physical world and agent S wants to have a successful answer to Q. Different criteria for what constitutes a "successful answer" lead to different positions in the realism vs. antirealism debate. Whereas Goldman proposes that a proposition B must be true in order to be a successful answer to the question, a constructive empiricist will settle for an empirically adequate answer. Remaining neutral between realism and antirealism, I will say that a proposition B is a solution/answer to a problem/question Q if and only if (i) B is a potential answer to Q and (ii) B satisfies the minimal criteria for success.

To remain neutral between truth vs. empirical adequacy, I will refer to the relevant constitutive aim of science as *planning adequacy*. Achieving planning adequacy would amount to being able to solve every scientific problem. Since the constitutive aim of other disciplines is presumably a form of planning adequacy as well, we should technically disambiguate *physical problem-solving adequacy* from other forms of planning adequacy (e.g. mathematical, logical, philosophical, etc.). For convenience, I will refer to 'physical problem-solving adequacy' as 'planning adequacy,' leaving it implicit that the focus here is on science. As I will argue in Section 4.5, neither truth nor empirical adequacy is actually sufficient for planning adequacy (at least not for beings like us). There is a further, subsidiary, constitutive aim necessary for achieving planning adequacy. Of course, the notion of 'adequacy' is ambiguous along various dimensions. What does it mean to 'really' solve a problem, and what does it mean to solve 'every' scientific problem? I turn to these issues next.

4.4 Ideal vs. Non-ideal Constitutive Aims

The minimal success criteria vary depending on whether we are talking about non-ideal vs. ideal scientific practice. There are at least two different dimensions of idealization. First, we can focus on obtaining exact—as opposed to approximate—solutions to scientific problems. This generates a distinction between exact vs. approximate empirical adequacy, which is simply a special case of the well-trod distinction between exact vs. approximate truth. Second, we can focus on all possible empirical problems, as opposed to the problems that scientific agents will actually face over their history. This distinction between solving all-possible vs. actual problems is a variety of modal completeness. Our knowledge about empirical problem-solving varies from radically incomplete at the beginning of science to maximally complete at the 'end of science.'

Both of these ideals—exactness and completeness—are *futuristic*: they will plausibly never be realized in finite time by any actual scientific community. What makes them futuristic is that they aim at "some ideally improved descendent that is never expected to actually exist, but that would result if the process of improvement went on forever"

(Lewis 1984, p. 230). In this way, *ideal* constitutive aims define the minimal criteria for success in the 'end of science.' In contrast, *non-ideal* constitutive aims define the minimal criteria for actual science, up to the demands that we require in practice and the problems that we actually face.[10]

The notion of 'adequacy' in 'empirical adequacy' is ambiguous along both dimensions of idealization. When it comes to solving actual problems non-ideally, it suffices for a solution to be approximate. Exact empirical adequacy (or truth) is unnecessary for practical purposes. As van Fraassen notes, "empirical adequacy is stronger than what counts as success *in practice*" (1995, pp. 157, 144). At best, scientists achieve various grades of approximate empirical adequacy, namely, an approximate fit between the predictions of their theory or solution procedure and the data obtained. At least due to measurement limitations, exactness is unattainable. Non-ideally, science aims at either approximate empirical adequacy or approximate truth. For this reason, scientific realism does not face a special problem of making sense of approximate truth that constructive empiricism avoids. Approximate empirical adequacy just is a special case of approximate truth: it is approximate truth about observables.

Restricted to individual problems, truth and empirical adequacy are both local criteria for success. They become global criteria when we consider wider classes of problems or problem-types. Beyond having (approximately) true or empirically adequate solutions to particular problems, scientists should ideally be able to solve any solvable problem they might encounter. This aim generates the second dimension of idealization: problem-solving completeness. Whereas non-ideal planning adequacy requires planning for all actual problems that will be encountered, ideal planning adequacy requires planning for all possible problems that could be encountered.

In the context of empirical problem-solving, the relevant notion of 'possible' is physically possible.[11] Whether a scientist ought to be able to solve a problem depends on whether that problem concerns a physically possible process, either exactly or approximately. Classical mechanics, for instance, typically concerns processes that are approximately physically possible (since in a non-classical world, an object can at best be ap-

[10] In passing, Davidson refers to a "constitutive ideal of rationality" (1980 [1970], p. 223). My contention here is that problem-solving adequacy is a constitutive ideal of science.

[11] Similarly, mathematics concerns at least logical possibility, while metaphysics focuses on metaphysically possible problem-solving. The local success criteria for solving a metaphysical problem are perhaps determined by the norms governing imagination.

proximately classical). Of course, we often do not know in advance which processes are physically possible. Hence, when inquiring, scientists often consider epistemic possibilities that turn out to be physically impossible. When scientists try to solve problems concerning non-physically possible scenarios, I take this to be instrumentally valuable for solving problems that are physically possible. Nonetheless, solving these non-physically possible problems is not a constitutive aim of science, even if this activity has independent epistemic value.

Here, I intend to remain neutral on how best to understand physical possibility. No doubt, different scientific realists have different preferred conceptions of physical possibility. Likewise, it is an open question which account of possibility is best suited for empiricism.[12] The specific content of scientific planning adequacy requires choosing an account of physical possibility, but the schematic conception remains the same.

My conception of planning adequacy is inspired by Gibbard's notion of a *hyperplan*.[13] Gibbard defines a hyperplan as a maximal contingency plan; it is a plan covering every possible circumstance (2003, p. 54). Clearly, no finite scientific agent ever needs a hyperplan to successfully solve the problems they actually encounter. Nevertheless, in the limit of successful scientific inquiry, we ought to develop hyperplans for all scientific problem-types. Even at the more local level of solving a given problem-type, scientists should ideally have a hyperplan for solving this problem-type.

To summarize: in the limit of a successful science, an epistemic community will have a plan covering all physically possible scenarios of interacting with the world. They will thereby attain *scientific planning adequacy*, i.e. knowledge of how to solve any (physically) possible physical problem. If we take empirical adequacy as the constitutive aim of problem-solving, then these agents will know all truths about observables, attaining what we might call *observational omniscience*. If instead we take truth as the minimal local success criterion, then agents in the end of science would know all physical truths and would thereby be *physically omniscient*. If a being attains physical omniscience, further scien-

[12]One strategy I pursue elsewhere is to combine an expressivist account of counterfactuals with Lange's (2009a) account of meta-laws, where counterfactuals are the truthmakers for physical necessity. This yields an empiricist-friendly account of meta-laws. I believe that the best way to interpret this account is as specifying the norms that govern our talk of laws and counterfactual planning, leading to a normativist approach to physical modality similar to Thomasson's (2020) approach to metaphysical modality.

[13]For his own inspiration, Gibbard cites Savage's (1954) discussion of big vs. small worlds in the context of decision theory.

tific inquiry becomes impossible. At most, they could 'inquire' in the sense of retrieving knowledge from their memory. But they could never inquire further about the physical world, since they already know everything about it.[14] As a result of their omniscience, they would have no need for scientific planning or problem-solving. All problems would already lay solved.

Considering the awesome power of a true scientific theory, one might worry that the notion of planning adequacy is redundant. Given either a true theory or an empirically adequate theory, wouldn't we *ipso facto* have achieved realist or empiricist planning adequacy, respectively? For recall that the notions of truth and empirical adequacy are futuristic: in the limit, they already purport to cover all past, present, and future states of affairs. Why then do we need a separate notion of planning adequacy? In the next section, I argue that for logically-omniscient beings, the notion of planning adequacy is genuinely redundant. However, for logically-imperfect beings like ourselves, neither empirical adequacy nor truth is sufficient. To achieve planning adequacy, beings like us must clarify the epistemic structure of our theories and problem-solving procedures. This, in turn, requires that we reformulate our scientific theories.

4.5 Aiming for Epistemic Suitability

Although an omniscient agent could not inquire about the world, mere *logical* omniscience does not preclude inquiry. Assuming that the state of the world does not follow from logic alone, a logically-omniscient being still has much to learn. If they were to arrive at an empirically adequate theory, then they would have reached the end of scientific inquiry, according to constructive empiricism. When coupled with logical omniscience, an empirically adequate theory automatically grants empiricist planning adequacy. Empirical adequacy entails that for any possible physical scenario, the theory has a model whose empirical substructure matches the observable phenomena. A logically-omniscient agent would immediately know how to apply such a theory to save the phenomena. As I argue below, this is because they already have full knowledge of the theory's epistemic

[14]As Woodard (forthcoming) argues, agents can genuinely check claims they already know. Hence, by 'omniscient,' I technically mean an agent who not only knows everything, but is also certain of this knowledge. Perhaps a physically omniscient agent could inquire further into the fundamental structure of reality, e.g. about how to state truths in a perfectly natural language. Whether these are further facts about the physical world depends on the relationship between metaphysics and science.

structure. Hence, for logically-omniscient beings, an empirically adequate theory is already sufficient for possessing a scientific hyperplan. Logically-omniscient agents with an empirically adequate theory are thereby guaranteed to be minimally successful when it comes to solving all possible scientific problems. They possess a plan that predicts and accurately describes all observable parts of the world. *Mutatis mutandis*, the same remarks apply to a logically-omniscient agent with a true theory. They are guaranteed to be minimally successful according to scientific realism.

However, most agents are presumably logically imperfect, rather than logically omniscient. Logically-imperfect agents have a real need to struggle through the derivations of what entails what, or of whether a given fact is necessary for another. They are not disposed to immediately know the logical connections between their concepts. For these agents, neither empirical adequacy nor truth is sufficient for planning adequacy. Instead, they have a further constitutive need to clarify the epistemic structure of their theories (insofar as they are interested in scientific problem-solving at all). To show that this need gives rise to a further constitutive aim, I will assume that these agents already have a true theory (so *a fortiori*, an empirically adequate theory as well).[15] This will be the case that humans find themselves in if we ever arrive at a true scientific theory, since we will never be logically omniscient (at least on account of our finite minds).

As argued in Section 4.3, a constitutive activity of science is empirical problem-solving. Successful problem-solving requires more than a true theory: it requires being able to apply the theory.[16] For logically-omniscient beings, truth alone (or empirical adequacy) suffices for applicability. This is not so for logically-imperfect beings. At one extreme, a true theory would be mostly useless for an agent who could not reason at all. To borrow an analogy from computability theory, having a true theory is akin to knowing a class of decidable problems. Agents who know this will know what they can and cannot do with their theory, in principle. However, this is a far cry from actually knowing a decision procedure, i.e. possessing a method for going from 'decidable' to 'decided.' Successful planning requires the latter, not merely the former.

To solve problems, agents need to figure out what suffices for a solution. This requires

[15]We could even assume that these agents know the logical relationships between their concepts but not that these are the logical relationships. Their thinking about these relationships could be (temporarily) caught in the musings of Lewis Carroll's tortoise.

[16]In a similar vein, Loewer notes that "the information in a theory needs to be extractable in a way that connects with the problems and matters that are of scientific interest" (2007, p. 325).

formulating the theory such that it provides a usable problem-solving procedure, such as an algorithm. For non-futuristic problem-solving, agents simply need to know a set of sufficient conditions for solving the problems they will actually face. In the ideal of futuristic problem-solving, agents need to be able to solve any possible problem they might encounter. This requires figuring out what they need to know to solve problems of any given problem-type. Overall then, scientific agents ought to determine what they need to know and what suffices to know in order to solve any possible scientific problem. These necessary and sufficient conditions for problem-solving are precisely what I have been calling *epistemic dependence relations* (EDRs). Hence, we see that clarifying the epistemic structure of a theory amounts to determining EDRs, i.e. epistemic relations relevant for problem-solving.

These considerations show that for logically-imperfect agents, there is a further constitutive aim of science, beyond truth or empirical adequacy. In order to be minimally successful, logically-imperfect agents must clarify the epistemic structure of their theories. Otherwise, they will be unable to achieve planning adequacy.[17] To clarify epistemic structure, they will need to determine the epistemic dependence relations for all possible problem-types. This will require reformulating their theories and attendant problem-solving procedures. In this way, reformulating is constitutive of the activity of clarifying epistemic structure.

Recognizing the need to clarify epistemic structure justifies the ubiquity of reformulations in science. Reformulations are not merely practically or instrumentally valuable tools for greater truth or empirical adequacy: they are essential for logically-imperfect scientific agents like ourselves. For such agents, reformulating is a constitutive aspect of scientific inquiry. By reformulating, we gain knowledge of a theory's epistemic structure. Reformulating thereby manifests the gaining of this knowledge, as opposed to merely putting us in a position to gain this knowledge. This contrast parallels a similar distinction that Sosa draws between constitutive vs. auxiliary intellectual virtues. On his view, some intellectual virtues are "knowledge-constitutive," as opposed to merely being auxiliary or practically or instrumentally useful for acquiring knowledge (2015, pp. 41–42).

For convenience, I will refer to the aim of clarifying epistemic structure as aiming

[17] My argument does not entail that logical omniscience is a further constitutive aim of science. Logical omniscience is not required for scientific planning adequacy. One does not need to know all logical truths to succeed at science.

for *epistemic suitability*. This terminology emphasizes that clarifying epistemic structure is necessary for having a suitable problem-solving procedure. It also places this aim within the same grammatical category as empirical adequacy. Just as both truth and empirical adequacy are formal features of theories, epistemic suitability admits a formal characterization. A theory formulation is *epistemically suitable* for solving problems of type P provided that it provides a problem-solving procedure for determining answers to P-problems. Some theory formulations are epistemically suitable for solving particular problem-types, but not others. In order to solve a particular problem using a theory, sometimes it is necessary to reformulate, leading to an alternative theory formulation. In this way, differences in epistemic suitability are analogous to different computer algorithms. One algorithm might be able to solve a problem that another cannot, despite the algorithms being in some sense aspects of "the same theory." To solve the problem using the latter algorithm, we must reformulate it, resulting in a new algorithm.[18]

Some aspects of scientific practice are instrumentally valuable: they are means to some further end. The ends themselves have *final value*: they are valuable in and of themselves, at least viewed from the lense of science.[19] I take it that the constitutive aims of science qualify as final ends of scientific inquiry. For instance, if science constitutively aims at truth, then truth acquires the status of being an end of science. Learning a truth would then have not only instrumental but also final value. In general, being a constitutive aim seems sufficient for being an end in itself (relative to that domain or activity).[20] Therefore, given that epistemic suitability is a constitutive aim, it too is an end of science. By constituting the achievement of this aim, reformulations thereby accrue a form of final value, as opposed to instrumental value alone. Gaining knowledge of epistemic dependence relations is non-instrumentally valuable, in the same way that making a theory more empirically adequate or true is non-instrumentally valuable. I take this to be the most promising argument against instrumentalism about reformulations.

[18]For an illuminating discussion of algorithms, see Goldman (2017, 22ff.). Although seeking testimony from experts *prima facie* counts as a solution procedure (e.g. writing a program that queries Wikipedia for answers), I view this as a subsidiary, practical aspect of scientific methodology. The ability to seek expert testimony is derivative on experts themselves having a non-testimonial solution procedure that articulates the relevant EDRs.

[19]For this distinction between final vs. instrumental value, see Korsgaard (1983).

[20]If science itself is only instrumentally valuable, then the ends of science are some further means to some further end. Settling foundational questions about what—if anything—possesses final value *tout court* lies outside the scope of my discussion here.

As a formal feature of theories, epistemic suitability is not a practical matter. At least, it is no more practical than truth or empirical adequacy. A theory formulation is either epistemically suitable to solve certain kinds of problems, or it is not. It either provides a solution procedure for these problems, or it does not. Nothing in this analysis depends intrinsically on practical features. Likewise, the constitutive aim of epistemic suitability is not a practical aim. At least, it is no more practical than the aims of truth or empirical adequacy. It is a component of the minimal success criteria for science. Insofar as science is an activity, these aims of course make reference to agents. Nevertheless, they do not rely on any special interests or idiosyncratic features of agents beyond their desire for doing science. In Section 4.7, I will contrast these constitutive aims with practical features of scientific inquiry, such as problem-solving speed. First, I consider and rebut objections that threaten to make epistemic suitability a pragmatic aim after all.

4.6 Is Epistemic Suitability a Pragmatic Aim?

Since the aim of epistemic suitability arises for logically-imperfect agents, one might worry that it is inherently pragmatic. This aim appears to arise from a feature of agents, namely their logical imperfection. Yet, the aims of truth or empirical adequacy are no less pragmatic in this regard.[21] As discussed in Section 4.3, the constitutive need for a true or empirically adequate theory also comes from features of most agents, namely that they are neither physically nor observationally omniscient (they have not yet achieved planning adequacy).

Rather than view these aims as arising for agents in particular, contingent epistemic situations, it is better to view them as universal aims of science. The constitutive aims of science apply to *any agent* engaged in science, no matter their epistemic situation. It is just that some agents already trivially satisfy certain constitutive aims. Observationally-omniscient agents trivially satisfy the aim of empirical adequacy. Logically-omniscient agents trivially satisfy the aim of epistemic suitability. Their commitment to science binds them to this aim, while particular features of their cognition conveniently discharge it.

[21]Indeed, one might worry that empirical adequacy *does* fundamentally appeal to agents, since what counts as 'observable' is relative to an epistemic community. Neither truth nor epistemic suitability depend on agents in this way. In response, a constructive empiricist could attempt to define observability in terms of the visual system of a computer.

Human scientists are bound to the same aims, without such cognitive advantages.

One might object that unlike these other aims, epistemic suitability references problem-solving, and problem-solving is an aspect of the *use* of a theory and to that extent pragmatic. Addressing this further worry requires distinguishing between (i) pragmatic (or practical) matters and (ii) (non-practical) epistemic matters. I will consider a few ways of drawing these contrasts. The upshot will be that if "pragmatic" is defined too broadly, then many matters of epistemic interest are pragmatic. Regardless, I take it that we can meaningfully distinguish between non-practical and practical epistemic matters (even if these matters are in some deeper sense 'pragmatic,' a claim that some pragmatists would endorse (Brandom 2011, p. 58)). It is at least in this sense that the aim of epistemic suitability is not pragmatic or practical, but distinctively intellectual (i.e. concerning non-practical, epistemic matters).

van Fraassen's Characterization

In developing constructive empiricism, van Fraassen posited a narrow conception of the epistemic and a broad conception of the pragmatic. According to van Fraassen, epistemic virtues bear on the truth or truth-conduciveness of a theory, concerning "how much belief is involved in theory acceptance" (1980, p. 4). In contrast, any virtue related to people counts as pragmatic, including how we apply theories. Specifically, van Fraassen says that "pragmatic virtues" concern "the use and usefulness of the theory; they provide reasons to prefer the theory independently of questions of truth" (1980, p. 88).[22] Likewise, he characterizes a "pragmatic factor" as "any factor which relates to the speaker or audience" (1980, p. 91).

On this construal of the pragmatic vs. the epistemic, the aim of epistemic suitability does seem to count as pragmatic (but so does much else, as we will see). Epistemic suitability concerns whether or not a theory formulation supports a problem-solving procedure for a given problem. Problem-solving is clearly an aspect of the use or usefulness of a scientific theory. At first glance then, constructive empiricism treats epistemic suitability as a pragmatic virtue.

However, matters are not so simple. In distinguishing the pragmatic from the epis-

[22]Earlier, van Fraasssen claims that "pragmatic virtues do not give us any reason over and above the evidence of the empirical data, for thinking that a theory is true" (1980, p. 4).

temic, van Fraassen is explicitly focused on the context of theory acceptance, where we are choosing between rival theories. This context is importantly different from that which arises when assessing compatible formulations. For in this context, we are not choosing between rivals. Instead, we are assessing whether we can gain knowledge of the world with a given theory formulation. If a particular theory formulation is not epistemically suitable for a given class of problems, then we cannot actually solve these problems using that formulation (we might need a different formulation). This results in a difference in what agents can gain knowledge about. If such differences in knowledge acquisition don't qualify as epistemic, then I don't know what does. It may be that differences in epistemic suitability do not always matter for non-pragmatically choosing between rival theories. Yet, there are other contexts of epistemic assessment besides this one.

Virtue Epistemology

Moreover, defining 'pragmatic' so broadly in all contexts leads to counterintuitive consequences. In particular, it implausibly restricts what counts as non-practically epistemic (i.e. intellectual). Standard assumptions in virtue epistemology provide a particularly dramatic illustration of this problem. In its various forms, virtue epistemology focuses on the intellectual virtues of believers, where these virtues can include abilities, dispositions, competences, or character traits. Here, a 'competence' can be understood as "a disposition to succeed in a given field of aimings" (Sosa 2015, p. 2). Knowledge is then understood as true belief arising from exercising one's intellectual virtue, e.g. an epistemic competence like apt perception. Notice that acquiring knowledge involves an agent to apply an ability. Consequently, van Fraassen's broad characterization of the 'pragmatic' seems to render all such epistemic acts as pragmatic.

Does this make virtue epistemology inherently pragmatic? Perhaps, but either way we can distinguish the intellectual from the practical even within virtue epistemology. We can meaningfully distinguish between virtues that seem constitutive of knowledge vs. virtues that seem merely practically or instrumentally valuable for acquiring knowledge. For instance, Sosa argues that a constitutive feature of some competences is that they manifest knowledge. Knowledge-constitutive competences are purely intellectual, as opposed to practical or auxiliary competences, such as being industrious rather than lazy when it comes to gathering evidence. This contrast underwrites a notion of "*purely* *intellectual* virtues, with no admixture of practical assessment" (Sosa 2015, p. 45).

The thrust of this response does not require adopting virtue epistemology. The point is rather that there are compelling ways to contrast epistemic and pragmatic matters that van Fraassen's 1980s characterization rules out in principle. I take this to indicate that we should not define 'pragmatic' so broadly, at least not when teasing apart practical issues from non-practical epistemic matters (even within a context that itself has pragmatic factors).

Muller's Characterization

Instead, I favor a broader notion of the epistemic and a narrower notion of the pragmatic, such as that proposed by Muller (2005) in his rendition of constructive empiricism. Personally, I would prefer to stop contrasting the epistemic with the pragmatic, and rather just talk about practical vs. non-practical dimensions of the epistemic. But the dominant words are already in use, so there is little I can do to alter settled conventions. Muller defines epistemic aspects of science as those features that are always relevant for deciding whether a proposition of an accepted theory counts as knowledge (2005, p. 63). Epistemic aspects include at least evidence, truth, and empirical adequacy. Unlike van Fraassen's characterization, Muller's entails that many aspects of problem-solving are genuinely epistemic. Figuring out what we need to know or what suffices to know to solve a problem clearly matters for deciding whether a possible solution is a genuine solution. In this sense, epistemic dependence relations are genuinely epistemic. Similarly, differences in epistemic suitability matter for deciding whether to accept a putative solution as knowledge.

In contrast, Muller defines *purely pragmatic* aspects as features that are *never* involved in deciding whether a proposition counts as knowledge. These include aspects of convenience, speed, and efficiency, considered in Section 4.7. More generally, pragmatic features are those that are *not always* relevant for deciding whether a proposition counts as knowledge.[23] On this narrower definition of the pragmatic, only some features of the use or usefulness of a theory are pragmatic. Again, since matters of epistemic suitability are involved in determining whether propositions count as knowledge, Muller's definition

[23] As Muller notes, constructive empiricists—but not scientific realists—view explanation and inference to the best explanation as pragmatic features.

properly classifies epistemic suitability as epistemic rather than pragmatic or practical.

In sum, the applicability of a theory has a distinctively epistemic, non-practical dimension. Problem-solving is not inherently pragmatic. We often use theories to provide reasons for believing or accepting claims about the way the world is. Theories provide these epistemic reasons largely by solving problems. Problem-solving is a central component of the epistemic role of theories. Moreover, the activity of problem-solving seems to be no more pragmatic than the activity of constructing a theory. Indeed, a common way of constructing theories is by unifying and connecting families of problem-solving procedures. If theory construction is not inherently pragmatic, then these constituent problem-solving procedures should not be inherently pragmatic either.

Still, one might worry that epistemic suitability simply tracks what an agent can or cannot do *in-practice*, where such considerations are inherently pragmatic. Given that an agent already has a true theory, can't they already solve *in-principle* any problem they might encounter? If their theory is futuristically true, it already in some sense "contains the answer" to any scientific question they might ask.[24] But truth alone (or empirical adequacy) does not entail that this solution is manifest or available, given your theory formulation. Assuming that epistemic reasons are *reasons for belief or acceptance*, the theory itself does not necessarily provide sufficient epistemic reasons to believe the answer it contains. To acquire these epistemic reasons, it is necessary to formulate the theory such that it is epistemically suitable for solving the given problem. This point is no more pragmatic than needing an (approximately) empirically adequate theory in order to have epistemic reasons in the first place.

To better understand the constitutive, non-pragmatic need for epistemic suitability, consider problem-solving contexts where reformulating is necessary. Perhaps your current theory formulation requires that you perform certain measurements in order to solve a given problem. What happens if you can't perform these measurements? This might be the case if your measurement device breaks, and—as a result of cosmically bad luck—remains broken forever. More prosaically, you simply might be set the problem: solve such-and-such problem without appealing to information that your current theory formulation requires. These are physically possible problem-solving contexts, and one ought

[24]Likewise, an empirically adequate theory is guaranteed to have a model whose empirical substructure is isomorphic to the relevant observable phenomena (van Fraassen 1980, p. 64).

to plan accordingly. Scientific activity involves not just planning for when things go well, but for when things go as poorly as possible. In these circumstances, reformulating your theory is *the only way* to make it applicable. This is not something that you merely have to do 'in-practice:' you have to do this in order to solve the given problem (or to know if it can still be solved at all). If instead you were logically omniscient, then you would already have at your disposal all non-trivially distinct theory formulations. There would be nothing further for you to do. But for logically-imperfect beings, there is a real, in-principle need to reformulate. Reformulations provide contingency plans: they offer alternative epistemic routes to a solution that although 'already encoded' in a true or empirically adequate theory, is not necessarily manifest or accessible.

4.7 A Need for Speed?

Are there any further constitutive aims of science, arising from further constitutive features of science? It is possible that there are. Indeed, my argument for an overarching constitutive aim of scientific planning adequacy—and its subsidiary aims—makes no claim to completeness. In particular, one might wonder whether agents of finite lifespan have a constitutive aim to solve problems as quickly as possible. Although one can pose purely theoretical questions about problem-solving speed or complexity, I will argue that these are separate from practical considerations that might motivate agents to solve problems quickly. Hence, solving problems as quickly as possible (or within other computational constraints) is never a constitutive aim of science, for anyone.

Consider scientific agents with finite computational resources, such as agents of finite lifespan. Is solving problems with minimal resources (or as quickly as possible) a constitutive aim for such agents? As an analogy, consider playing chess within finite time (as most chess is played). Given that a chess player wants to win, they have a constitutive aim to try to checkmate before their time expires. Likewise, given that a human scientist wants to solve a scientific problem, they have a constitutive aim to try to solve it within their lifetime. In both cases, these are constitutive aims *at the agent-level.* As discussed in Section 4.2, both realism and empiricism focus on constitutive aims *at the activity-level*, i.e. the constitutive aims of Chess or Science. Constitutive aims at the agent-level do not entail constitutive aims at the activity-level. A feature or success criterion can be constitutive for an agent without being constitutive of the activity at large.

Considered anthropomorphically, the game of Chess does not care which player wins (or if any player wins at all). Likewise, Science does not care which scientist solves any particular problem, or when. Hence, the constitutive need for a finite agent to try to solve problems as quickly as possible does not entail a constitutive need for speed at the activity-level. As further support for this claim, recall that the ideal constitutive aims of science are futuristic: the minimal criteria for ideal success concern an infinite time frame, with limitless computational resources. Within this futuristic idealization, we treat the activity as having an arbitrarily-long time frame. Even though particular games of chess or scientific problem attempts are finite, the futuristic, ideal aims abstract away from temporal constraints. The constitutive aim of science remains planning adequacy. The computational limits of individual scientists do not affect this constitutive aim, nor do they create an additional constitutive aim.

Although maximizing problem-solving speed or minimizing computational resources are not constitutive aims of science, planning adequacy already encompass related theoretical questions about speed or resource minimization. It *is* constitutive of science to seek knowing how quickly a problem can be solved, along with knowledge of a procedure for solving it that quickly. Questions of speed or resource consumption are themselves theoretical questions, such as *what is the fastest physically possible solution to this problem?* These theoretical questions correspond to problem-solving contexts where the problem is to solve a given sub-problem within a certain period of time. Since scientists ought to prepare for any possible problem-solving context, they ought to investigate the speed at which they can solve problems. The question of how quickly a problem can be solved is yet another scientific question, generating its own problem to solve. These investigations are a purely theoretical aspect of science (and mathematics), with attendant literatures, e.g. on the minimization of proof length or the limits of computer speed. Scientists ought to investigate questions about speed and computational resources, regardless of any practical benefits such investigations might yield.

Lest there remain a whiff of paradox, it is vital to distinguish two different kinds of goals: (i) wanting to know how quickly a problem can be solved vs. (ii) wanting to solve a problem as quickly as possible. The former is a purely theoretical aim (and includes wanting knowledge of how to solve a problem as quickly as possible—e.g. if the world demands it). The latter is a genuinely practical or pragmatic aim. Science itself does not require one to care about how quickly they will actually solve problems (as opposed to how quickly a problem could be solved). Of course, the motivations for asking questions about speed are often practical, borne of a desire to have the fruits of an investigation sooner rather than later. Nonetheless, we can do science as slowly or quickly as we please, while doing science all the same. As Peirce notes, it is constitutive of inquiry to hope that there is an answer to our problem.[25] It is not constitutive (of collective inquiry) to hope that you yourself will find that answer within your lifetime, although you very well might hope this.

One might have thought that desiring speed is non-constitutive because speed is species-relative. What counts as fast for one kind of finite being might be hopelessly slow for another. Yet the empiricist notion of observability is already species-relative, as are the various grades of logical imperfection. Due to species-relativity, the content of the aim of empirical adequacy is necessarily indexical. So it is not this indexical nature which makes a desire for speed into a practical aim, preventing it from being a constitutive aim of science. Instead, a desire for speed is non-constitutive because it arises from the extra-scientific preferences of particular agents, relative to their particular constraints. The only preference that matters for the constitutive aims of science is the desire to do science. Some agents might aim to solve scientific problems quickly or within a reasonable time, but these are practical goals. They typically arise from contingent social structures, such as temporal requirements for career advancement. A finite scientific agent will labor under more practical constraints than an infinite one, but their constitutive aims remain the same. The need for speed is at most a non-constitutive practical aim, at the agentive level.

4.8 Why Talk about Aims at All?

Some may be skeptical that it makes sense to talk about science as having aims. Such worries might stem from a more general skepticism toward 'the aims of activities,' such as

[25]Misak quotes Peirce as saying that "the only assumption upon which [we] can act rationally is the hope of success," which she interprets as meaning that "it is a regulative assumption of inquiry that, for any matter into which we are inquiring, we would find an answer to the question that is pressing in on us" (2013, p. 50).

belief, inquiry, or assertion. Considering the so-called aim(s) of belief, Wedgwood rightly quips that "Beliefs are not little archers armed with little bows and arrows: they do not literally 'aim' at anything" (2002, p. 267). The same can of course be said for science: if it makes sense to talk collectively about science at all, science is not literally aiming at anything. Nonetheless, Wedgwood goes on to give a normative explication of claims about the aim(s) of belief: such claims state correctness conditions for beliefs. Here, I have pursued a similar strategy: constitutive aims of science provide criteria for minimal success in scientific endeavors.[26]

A common objection to aims-talk is that it presupposes a univocal or hegemonic aim, where none is to be found. In his earlier work, Putnam expressed this criticism in a couple places, remarking that "it is hard to believe that there is such a thing as 'the aim of science'—there are many aims of many scientists" (1979 [1971], p. 355).[27] Earlier, he remarked that "The use of such expressions as 'the aim of science'…is already extremely apt to be misleading. For there is no *one* 'aim of science'….Different scientists have different purposes" (1979 [1965], p. 233). More recently, MacFarlane has expressed a similar worry in the context of the aim of assertions. Considering the idea that assertions uniquely aim at the truth, MacFarlane says, "This idea is pretty obscure anyway. Even if truth is *an* internal normative aim of assertion, it is certainly not the only such aim" (2005, p. 227).

In developing my account of constitutive scientific aims, I have not assumed that science has a unique aim. Instead, I have focused on identifying a constitutive aim of science, namely, planning adequacy. I happily allow that science might have other constitutive aims and of course many other non-constitutive aims. As a practical matter, scientists presumably are trying to solve their most pressing problems as quickly as their energy and resources allow.

Part of Putnam's misgivings arise from the fact that individual scientists can themselves have many aims. I take this part of Putnam's worry to lose its force once we distinguish between aims at the agential-level vs. at the activity-level, as discussed in Section 4.2. As I hope to have made clear, constitutive aims of science are largely insulated from the aims of individual scientists. In order for there to be constitutive aims of

[26]More broadly, Sosa argues that the notion of a constitutive aim is useful for understanding action, perception, and knowledge: "We find unity across action, perception, and knowledge. All three are constituted by aimings, by performances with a constitutive aim" (2015, p. 24).

[27]I thank Gordon Belot for suggesting I review this essay by Putnam.

science, it suffices that there are scientific agents aiming to do science.[28]

The primary reason to talk about the aims of science is so simple that it is worth reiterating. In order to determine what it means for a scientific activity to succeed, we must talk in terms of aims, goals, or purposes. These aims or goals define success criteria. As van Fraassen notes, disagreements about "the aim pursued in science" lead to disagreements about what "counts as scientific success" (1995, p. 143). Constitutive aims play a special role in characterizing minimal success criteria for an activity. In this vein, Putnam himself invokes "the aims of inquiry" in order to criticize conventionalist approaches to logic, his point being that we need assurances that such approaches do not interfere with the aim of having "a true description of the world" (1979 [1968], p. 188).

If an activity like science does not have criteria for minimal success, then it loses much of its intellectual interest. If science doesn't aim at anything, then scientists might as well be flailing their arms around. And what intellectual grounds would we have to criticize them? Insofar as we do have grounds to criticize methodologically-wayward scientists, I take there to be a meaningful notion of constitutive aims for science. When we talk about scientists who are failing to do science, we presuppose that there is at least some overarching aim(s) to which such scientists are failing to contribute. Such scientists are not even contributing to the minimal success of science.

Hence, I do not see talk of aims as being any more problematic than talk of correctness conditions, or success conditions more generally. Insofar as particular practices have standards for success, those practices have aims. There are many philosophical concepts that I find deeply mysterious. Aims-talk is not one of them.

4.9 Conclusion

I have argued that one constitutive aim of science is to solve (all possible) problems about the physical world. As a constitutive aim, this specifies a minimal success criterion for science. Scientific realists contend that minimal success requires a true scientific theory. Antirealists contend that truth is not necessary for minimal success. In particular, constructive empiricists argue that empirically adequate solutions to scientific problems

[28]Plausibly, we do not even need this much. We can talk about the constitutive aims of chess even if no one plays chess ever again. Constitutive aims characterize criteria for minimal success in an activity *were there to be agents undertaking said activity.*

would suffice. Since resolving this contentious issue is unnecessary for my larger argument, I have remained neutral on the specific minimal criteria for solving a scientific problem.

Instead, I have argued that—as presently formulated—both realism and antirealism neglect a further constitutive aim of science. Planning adequacy requires more than an empirically adequate or even true theory. In order to solve problems, we require an epistemically suitable formulation of a relevant theory or problem-solving procedure. The aim of epistemic suitability requires that we identify a sufficient knowledge-base for reaching a solution. Ideally, we should identify what we need to know in order to solve scientific problems. Determining what we need to know or what suffices to know to solve problems requires that scientists reformulate their theories and problem-solving procedures. Reformulations are thereby essential for attaining epistemic suitability. For this reason, reformulations are non-instrumentally valuable. They are valuable as scientific ends in themselves, not merely in the service of greater truth or empirical adequacy.

Each constitutive aim has both ideal and non-ideal dimensions. The non-ideal constitutive aims of science concern what science has to achieve for minimal success within its actual history. Non-ideally, we require epistemically suitable formulations only for the problems that we will actually face. Solving problems non-ideally requires only approximate truth or approximate empirical adequacy, where our actual practices settle what counts as 'good enough.' The ideal constitutive aims of science specify what science would have to achieve for minimal success in the limit of infinite time and resources. Ideally, science would be able to solve any physically possible problem about the physical world. This would require complete epistemic suitability for all physical problem-types. Additionally, ideal solutions would be as true or empirically adequate as physically possible, e.g. up to the limits of physically possible measurement precision.

We can distinguish the constitutive aims from additional non-constitutive aims. No doubt, science would go better if we developed faster problem-solving procedures. No doubt, it would be a great success if we arrived at a fundamental language for describing reality. Although these achievements would constitute scientific success, they are arguably not minimal success criteria for science.[29] They can be ideal aims, without being

[29]Fundamentalist scientific realists might argue that seeking a canonical language for reality *is* a constitutive aim of science. I take the arguments in Section 1.5 and Chapter 5 to undermine fundamentalism.

ideal constitutive aims. If constructive empiricism is correct, then truth is also a non-constitutive aim, along with providing explanations of physical phenomena. My defense of epistemic suitability advantageously accounts for the value of reformulating without appealing to explanations or differences in explanatory goodness. It thereby provides an account that realists and antirealists alike can endorse.

Joshua Robert Hunt

joshhunt@umich.edu

ORCID iD: 0000-0001-5150-0388

Chapter 4 **Reformulation as a Constitutive Aim of Science** **127**
4.1 Introduction 127
4.2 Constitutive Aims 128
4.3 Scientific Planning as Problem-Solving 133
4.4 Ideal vs. Non-ideal Constitutive Aims 134
4.5 Aiming for Epistemic Suitability 137
4.6 Is Epistemic Suitability a Pragmatic Aim? 141
4.7 A Need for Speed? 146
4.8 Why Talk about Aims at All? 148
4.9 Conclusion 150

Chapter 5 **Making it Manifest: The Intellectual Value of Good Variables** **153**
5.1 Hidden Symmetries and Manifest Properties 153
5.2 Manifest vs. Hidden Facts 155
5.2.1 Simple examples, on the sleeves 159
5.3 Manifest Meanings 164
5.4 The Value of Making it Manifest 166
5.4.1 Ruling out epistemically possible solutions 168
5.4.2 Less surprising, more intelligible 171
5.4.3 Degrees of manifestness 172
5.4.4 Problem-solving adequacy and fruitfulness 173
5.5 Coordinate Transformations 174
5.6 Preferences, Fundamentality, and Privileging 180
5.7 Gauge Choices 185
5.8 Manifest vs. Hidden Symmetries of Hydrogen 189
5.9 Hidden Symmetries in $\mathcal{N} = 4$ super Yang–Mills Theory 195
5.10 Conclusion 203

References **205**