

This is the penultimate version (prior to copyediting) of chapter 1 of
A. Müller, *Constructing Practical Reasons*, Oxford: Oxford University Press, 2020.
Please cite the published version, which is accessible via Oxford Scholarship Online:
<https://doi.org/10.1093/oso/9780198754329.001.0001>.

Chapter 1

What is constructivism?

1.1 The basic idea

Constructivism, as I understand it here, is first and foremost a view about the relation between thought and its object. To be a constructivist about a certain domain is to take the facts of that domain to be an upshot, rather than the measure, of correct judgement in that domain. A constructivist about practical reasons thus denies that, when an agent comes to realize that something is a reason for a certain action, she does so *in response to* the corresponding normative fact, that is, a fact about what is a reason for doing what. Instead, such facts are somehow explicable in terms of our reason judgements. That does not mean that the facts are whatever we think they are. Constructivists do not deny that we can be mistaken about what is a reason and what is not. But they do suggest that normative facts are ultimately grounded in our mental states or activities. In particular, according to the constructivist account of practical reasons that I will present in this book, facts about what is a reason for doing what obtain in virtue of the activity of practical reasoning and the role that the corresponding reason judgements play in it. More specifically, such a fact obtains because the corresponding reason judgement is true, and the judgement is true because the episode of reasoning that it is apt to guide is sound.

In this first chapter, I want to start by presenting the main aspects of the constructivist account I just sketched in a bit more detail. I do this, firstly, in order to give you an overview and to motivate the remaining chapters of this book, where these aspects and the questions to which they give rise will be discussed at much greater length. But I also do this because the term “constructivism” has been used to pick out a variety of rather

different positions in metaethics, both by advocates and by opponents of the views so labelled. Hence, I want to be as clear as possible about how and why I use that term, to avoid adding to the confusion. Let me be clear, though, that I do not claim that my proposed understanding of constructivism somehow synthesizes or subsumes all of the various conceptions of constructivism that can be found in the literature.¹ The view that I put forward adopts many important insights from those other conceptions (as I will point out along the way). In fact, little of what I have to say in what follows has never been suggested before. But my aim here is not to identify a common denominator; I am sceptical that one could be found. The basic idea of constructivism is essentially a metaphor: some things are *constructed* from certain other things; they are *made*, rather than discovered, as the well-known slogan has it. Such metaphors can be spelled out in different ways, none of which is, as such, the right one. Nevertheless, some ways might be more useful or more interesting than others. The position that I present in this book is meant to be true to the idea encapsulated in that metaphor, and it takes up some important thoughts of paradigmatic constructivists like John Rawls and Christine Korsgaard, as we will see in the following sections.² But most importantly, it is meant to offer a distinct and *prima facie* attractive account of practical reasons.

1.2 Constructivism and representation

Let me now elaborate on the main aspects of the constructivist position that I sketched above. I will start with its account of the nature of reason judgements in this section, and then address its understanding of truth as well as the mind-dependent status it assigns to reason facts in sections 1.3 and 1.4.

¹ Advocates of constructivist views in metaethics include John Rawls (1980; 1993), Christine Korsgaard (1996; 2003; 2009a; 2009b), Carla Bagnoli (2002; 2012; 2013b), Aaron James (2007; 2012) and Sharon Street (2008; 2010; 2012). For a comprehensive overview, see Bagnoli (2017).

² Constructivism is often considered to have a Kantian pedigree, and both Rawls and Korsgaard's views are certainly inspired by Kant. However, whether Kant is best interpreted as a constructivist is controversial among Kant scholars, and I will not address this issue here. For a helpful discussion of metaethical constructivism that pays special attention to its roots in Kant's philosophy, see Schafer (2015a; 2015b).

Constructivism rejects a familiar conception of the kind of attitude that is involved in thinking about reasons. Consider the following example. Lara and her friends are out for dinner. They are about to order from the menu, which, this being a very small restaurant, offers only two options for the main course: a stuffed squash and duck breast marinated in hoisin sauce. Lara is unsure what to choose, so she confers with her friends, who have been here before. They tell her that the squash's stuffing is vegetarian, and that the duck breast comes with rice. Lara also learns that hoisin sauce contains soybeans, and this she takes to be a reason against having the duck breast (she is allergic to soybeans). This reason judgement settles her choice, and she proceeds to order the squash. What role did Lara's judgement that the fact that hoisin sauce contains soybeans is a reason against ordering the duck breast play in the process that lead to her decision? What kind of attitude do we describe an agent as having when we attribute such a reason judgement to her?

The familiar conception answers those questions like this: To believe that the duck breast comes with rice and that hoisin sauce contains soybeans is to represent certain facts as obtaining, and when those facts really do obtain, as we shall assume, then those representations are accurate. Hence, by forming those beliefs, Lara acquired new information about the options in a choice problem she currently faces, information that she can apply in making a decision. The familiar conception extends this account of the cognitive role of ordinary beliefs to the case of reason judgements. Her reason judgement,³ if true, also provides Lara with information about the options in her choice problem, information that will help her to come to a decision. Of course, the beliefs and the reason judgements differ in the kind of information they provide. The former

³ I refer to these attitudes as "reason judgements" rather than, e.g., "beliefs about reasons", because the latter is often associated with a representationalist account of the function of the attitudes in question. As will become clear shortly, constructivism rejects such an account of the attitudes' cognitive function. So, in order to avoid prejudging the issue, I am using what I take to be a less committing term. "Judgement" is thus meant to pick out a certain kind of mental state, not the event of forming such a state.

provide non-normative information, that is, they accurately represent certain non-normative facts about the options in question, such as the fact that the duck breast comes with rice, or that hoisin sauce contains soybeans. To form a reason judgement, on the other hand, is to acquire normative information about the choice situation and to represent some normative fact, some fact about what is a reason for what, as obtaining. But underlying such differences in content—in what kinds of facts they are about—is a fundamental similarity in the role those two attitudes play in the cognitive economy of an agent who faces a choice. Just like ordinary beliefs about non-normative matters, reason judgements are representational states which, if accurate, provide the subject with information that she can apply in making a decision.

The details of this conception can be filled in in numerous different ways, particularly when it comes to characterizing the facts that reason judgements purport to represent. Are they *sui generis* normative facts, or can they somehow be reduced to other, non-normative facts? Are they causally efficacious? How can we gain knowledge of them? These are just some of the questions on which different version of this conception can disagree, and they are certainly important questions. Nevertheless, we can put them to one side for the moment, because the fundamental aspects of the conception, the ones that will be most relevant in what follows, are not affected by how these questions are answered. These aspects concern the relation between those reason facts on the one hand and our reason judgements on the other, as well as the role reason judgements play in the cognitive processes of agents like Lara. According to the familiar conception, judgements about what is a reason for what relate to certain facts in the same way that the belief that hoisin sauce contains soybeans relates to the fact that hoisin sauce contains soybeans: they purport to represent those facts. Hence, they provide an agent with information that is relevant to her choice, if they represent accurately.

This conception of the nature of reason judgements is quite popular in contemporary metaethics. The idea that reason judgements purport to represent certain facts and, if accurate, provide agents with information about the normative aspects of her situation seems to underlie, in one of its guises, the positions developed by David Enoch (2007; 2011a), Derek Parfit (2011b), Thomas Scanlon (2003; 2014), Mark Schroeder (2007),

Michael Smith (1994), and many others. It is often not so much articulated explicitly as hidden in the assumption that reason judgements and ordinary beliefs about non-normative matters are the same kind of attitude towards different contents. If you take the main difference between the two to consist in the facts they represent, it is understandable why you will focus on giving an account of the facts our reason judgements are about, while saying little about the kind of attitude you take those judgements to be—there is no need to, if it is the same as our ordinary beliefs.

Constructivism offers an alternative to this familiar conception. It suggests that the cognitive function of reason judgements differs fundamentally from the cognitive function of ordinary beliefs about non-normative matters. According to the constructivist, reason judgements play a distinct role in the psychology of a deliberating agent: the point of those judgements is not to provide an agent with additional information that she can apply in making her decision. Instead, reason judgements *guide* her in making a decision on the basis of the information she has. To use a computational analogy: while ordinary beliefs can provide the input for decision-making processes, reason judgements provide the instructions that such processes follow in translating those inputs into outputs, that is, into decisions. Thus, when Lara forms the belief that the hoisin sauce contains soybeans, she acquires new information about one of her options, but when she forms the judgement that this is a reason not to order the duck breast, she does not acquire more information. Instead, this judgement will allow her to apply the information she already has to the choice she faces. Information is not enough to come to a decision; we also need to have an idea of how to bring any such information to bear on the problem of choosing one of the available options. This is the distinct contribution reason judgements make to the process of reasoned decision-making.

To be more specific about the main point of disagreement between constructivism and the familiar conception of the nature of reason judgements outlined above, let me introduce the following thesis about what kind of attitude reason judgements are:

Representationalism

Reason judgements are representational states. Their cognitive function is to represent certain facts, and if they do so accurately, they provide the agent with normative information by indicating that a certain fact concerning what is a reason for what obtains.

Proponents of the familiar conception accept Representationalism, or one particular version of it, while constructivists reject it altogether and take reason judgements to play a fundamentally different cognitive role than, for example, beliefs about the composition of some condiment. Constructivism is thus a form of anti-representationalism about the nature of reason judgements.⁴

Constructivists are not the only ones who reject Representationalism. They share this negative part of their view with expressivists such as Simon Blackburn (1998) and Allan Gibbard (1990; 2003). The two views differ, however, in their positive accounts of the kind of mental states they take normative judgements to be. Expressivists typically conceive of them as motivational or desire-like states, that is, states that are characterized in terms of their role in the explanation of an agent's actions. Constructivists, on the other hand, think that reason judgements are a distinct kind of mental state, the nature of which can only be understood by attending to its role in guiding the agent's practical reasoning. In other words, whereas expressivism characterizes reason judgements as a kind of pro-attitude, constructivism identifies them in terms of their role in a process that leads to the formation of such attitudes. The two views thus disagree substantially on the cognitive role of the attitudes in question.⁵

⁴ Carla Bagnoli (2002, 125) agrees that constructivists do not consider moral judgement to involve the "aspiration to correctly represent a moral reality".

⁵ Sharon Street (2008; 2010; 2012) also argues that the attitudes that provide the basis for her constructivist account of practical reasons must be distinguished from ordinary desires. But she nevertheless characterizes them as attitudes that are "by their nature motivating" (2008, 230), which seems to make her account of normative judgements more difficult to distinguish from the expressivist's than the reasoning-based account that I propose in this book.

Representationalism is also rejected by global expressivists such as Robert Brandom (1994) and Huw Price (2011; 2013). Moreover, Brandom's claim that the function of normative vocabulary is to make explicit "an attitude that otherwise could be implicit only in what is done—namely, the endorsement of a pattern of practical reasoning" (1994, 271) seems to attribute to reason judgements a cognitive role similar to the one described in chapter 2 below. The constructivist theory presented in this book indeed bears important similarities to the views of Brandom and Price. But there are also significant differences. For one, constructivism about practical reasons is a *local* theory that contrasts the cognitive function of reason judgements to that of, e.g., beliefs about one's environment—it does not reject Representationalism across the board. Constructivists also do not share Brandom's and Price's commitment to deflationism (cf. chapter 4).

John Rawls, probably the most influential proponent of constructivism in the twentieth century and the author responsible for introducing that label into the metaethical debate,⁶ agrees that constructivists reject the idea that normative judgements resemble ordinary beliefs in playing a representational role in an agent's psychology. He contrasts constructivism with a view he calls "rational intuitionism."⁷ This view holds that "moral first principles and judgements, when correct, are true statements about an independent order of moral values" and they are "known by theoretical reason", which means that their epistemology is structurally similar to that of empirical judgements: "moral knowledge is gained in part by a kind of perception and intuition" (1993, 91–92). The

⁶ It is important to note that Rawls's views on constructivism have changed over time, or at least shifted in focus. Whereas his Dewey lectures (1980) are simply concerned with "Kantian constructivism", he later distinguishes between "political constructivism" and "Kant's moral constructivism" (1993, 89–129). Only the latter of these two is what he calls a "comprehensive moral doctrine", that is, a view that also addresses metaphysical and epistemological questions and thus qualifies as a metaethical view. Hence, I will focus on Rawls' characterization of constructivism as such a comprehensive doctrine.

⁷ Cf. Rawls (1980, 557–64; 1993, 91–101). He identifies Samuel Clarke, Richard Price, Henry Sidgwick, G. E. Moore and David Ross as paradigmatic advocates of rational intuitionism.

rational intuitionist's conception of the relation between moral judgements and moral facts is thus very similar to Representationalism. According to Rawls, constructivism rejects this conception, along with the "sparse conception of the person as [...] a knower" (1993, 92) that goes with it. Rawls's characterization of constructivism thus puts the view in fundamental opposition to Representationalism.

Christine Korsgaard is the second major proponent of constructivism in the last decades. Especially in her more recent writings, she puts the contrast between constructivism and realism, which she considers to be its main competitor, in terms of the function or purpose of moral and other normative concepts:

[T]he difference between a realist and a constructivist theory rests in the way the two views understand the function of concepts [...]. A realist believes that the function of concepts is to describe the world, to mark out the entities we find there, while a constructivist believes that the function of (at least some) concepts is to mark out, in a schematic way, the solution to some problem that we face. (2008, 22)

Constructivism thus rejects the idea that the cognitive function of normative judgements is to represent reality. That is not what normative concepts are for, their "cognitive job" is not a descriptive one (2003, 105). Korsgaard thus clearly rejects a representationalist account of normative thought. Like Rawls, she understands constructivism as a view that is meant to offer an alternative to Representationalism.⁸

1.3 Constructivism and truth

Constructivism offers an explanation of why some things are reasons for certain actions. It explains such reason facts, in the first instance, in terms of the truth of the

⁸ Her alternative proposal for the function of normative concepts is admittedly less clear. On her preferred way of formulating that alternative, those concepts denote, in a "formal" or "schematic" way, the solution to some problem we face insofar as we are agents (1996, 113). While this does not strike me as being in conflict with the view that I develop in this book, I am not confident enough in my understanding of her proposal to make the case for any substantive agreement.

corresponding reason judgement: they obtain because that judgement is true.⁹

Constructivists thus do not deny that there are facts about what is a reason for doing what. Nor do they deny that those facts correspond to true reason judgements, as long as that claim is understood as a simple biconditional:

Correspondence

The judgement that p is true if and only if it is a fact that p.

This claim is symmetrical and thus leaves open the question of explanatory priority between its left- and right-hand sides. It is on this latter issue that constructivists disagree with their opponents. They maintain that the order of determination in Correspondence goes from left to right: they treat reason facts as the explanandum rather than the explanans.¹⁰ What motivates this reversal is the constructivist's rejection of Representationalism as an account of the cognitive function of reason judgements. Truth, it seems, is closely connected to successful functioning. So, if the cognitive function of reason judgements is not to represent certain facts, then their truth will not be grounded in the accurate representation of such facts, either. Denying that reason judgements are true, if they are true, because they accurately represent certain reason facts, allows constructivists to employ Correspondence as part of an explanation of its right-hand side, rather than its left-hand side. For them, a reason fact is not the ground of a reason judgement's truth, but rather its ontological shadow.¹¹

⁹ Note that what explains a reason fact according to this proposal is the *truth* of the corresponding reason judgement, not the true judgement. That is, no actual judgement—understood as some specific psychological reality—is necessary for a reason fact to obtain. What is required, rather, is that the conditions that ground the truth of that judgement are actually met.

¹⁰ Compare Bagnoli (2013b, 167–68), who writes that constructivists reject the claim “that moral judgments are truth-evaluable because they correspond to or represent some portion of reality” and that, rather than holding “that [moral] facts are truth-makers,” they hold “that such a domain [of facts; A.M] is the result of the activity of practical reasoning”.

¹¹ This way of drawing the contrast between constructivism and its alternatives takes cues from Kit Fine's distinction between “antifactualist” and “factualist” accounts of a certain practice:

For their explanatory project to succeed, constructivists must offer an alternative account of what it is for reason judgements to be true. Here, the threat of circularity looms. To avoid it, the constructivist's account of the truth of reason judgements must not presuppose the prior existence of any reason facts. Moreover, their account must understand truth in terms that are substantive enough to bear the explanatory weight that the constructivist's use of Correspondence puts on them. Providing such an account is one of the main challenges that advocates of constructivism have to meet in order to deliver on the promise that their theory of practical reasons holds. As I will argue in chapters 4 and 5, the best way to meet this challenge is for constructivists to build on their account of the nature of reason judgements as an attitude that guides practical reasoning and propose that the truth of those judgements is to be understood in terms of the soundness of the episodes of reasoning that they are apt to guide.¹²

Offering a constructivist account of some domain thus has consequences for understanding what the truth of judgements in that domain amounts to. Here, too, I am in agreement with how both Rawls and Korsgaard characterize constructivism. Rawls

The factualist's account must in this sense be representational: it must link up the practice with the underlying facts or subject matter, while the antifactualist's account will be nonrepresentational. In the one case, the practice must be seen as engaging with the possible facts and it must be understood—at least, in part—in terms of how it engages with those facts. In the other case, the practice is taken to be disengaged from the facts; and rather than understanding the practice in terms of how it represents the possible facts, the facts themselves should be understood in terms of how they are 'projected' by the practice. (Fine 2001, 24)

Note, however, that despite the similarities between “shadows” and “projections”, constructivism differs importantly from Simon Blackburn's expressivist projectivism, which has projected facts correspond to whatever normative judgement an agent actually endorses; see Blackburn (1988b) and the introduction to Blackburn (1993).

¹² Showing how a constructivist account of practical reasons can avoid circularity is also a major concern in James (2007). James argues that this commits constructivists to showing that the norms of practical reasoning do not derive their authority from the fact that they lead to reason judgements that are true on independent grounds (2007, 308). The account of correct practical reasoning presented in chapter 5 allows constructivists to discharge that commitment.

acknowledges that because they reject the rational intuitionists' account of moral judgements, constructivists cannot "conceiv[e] of truth in a traditional way by viewing moral judgements as true when they are both about and accurate to the independent order of moral values" (1993, 92). In other words, constructivists cannot understand truth in terms of accurate representation. Rawls seems unsure what to make of this consequence. Sometimes he recommends that constructivists should not think of moral judgements or principles as being true but rather as being the most reasonable ones, highlighting that this allows us to avoid any commitments to a theory of truth that deviates from the standard representationalist account that is most plausible in the case of ordinary empirical judgements (1980, 554). Other times, however, Rawls seems to suggest that constructivism puts forward an alternative conception of the truth of moral principles and judgements, one that understands their truth in terms of the notion of reasonableness and which allows constructivists to affirm the genuine truth-aptness of those judgements, despite their rejection of the intuitionist's representationalist account of them (1980, 569; 1993, 126n).

Korsgaard is less ambivalent when it comes to the constructivist's commitments in the theory of truth. Her constructivism and what she calls substantive realism do not disagree on whether or not moral judgements can be true, but on what makes them true, if they are.¹³ The substantive realist claims that there are correct answers to moral questions "*because there are moral facts [...] that those questions ask about*" (1996, 35; her italics).¹⁴ In other words, an answer to a moral question is correct because it gets the

¹³ Hence, they are both versions of realism in a broader sense, which includes anyone who believes that there are correct answers to moral (and other normative) questions and thus thinks that "ethics isn't hopeless" (1996, 34–35). Korsgaard calls this broad version "procedural realism", presumably because she thinks that the availability of procedures that will lead us to correct answers to moral questions is a necessary condition, if not for there being such answers, then at least for the pursuit of those answers to not be a hopeless endeavour.

¹⁴ Korsgaard counts Samuel Clarke, Richard Price, G. E. Moore, David Ross, H. A. Pritchard, Peter Railton, David Brink, Thomas Nagel, and Derek Parfit, among others, as proponents of substantive realism (1996, 19; 2003, 101). Note the similarity between this list and the list of philosophers Rawls identifies as "rational intuitionists" (see note 7).

moral facts right. Korsgaard's constructivism, by contrast, is meant to give an account of what makes answers to moral questions correct that does not refer to any antecedent moral facts. She rejects the idea that the procedures which lead us to correct answers to moral questions are procedures that "track" certain "facts which exist independently of those procedures" and thereby provide "ways of finding out about a certain part of the world, the normative part" (1996, 36–37).¹⁵ This does not prompt her to be sceptical about the truth-aptness of moral judgements, though. All that is required for truth, Korsgaard claims, is that there is a standard for the correct use of the relevant concepts: "when a concept is applied correctly, what we get is truth" (2003, 117). Such a standard is provided by what she calls "the correct conception" of the concept in question. But what makes a conception correct, in the case of moral and other normative concepts, is "that it solves the problem,^[16] not that it describes some piece of external reality." (2003, 117)

1.4 Constructivism and mind-dependence

According to constructivism, that some things are reasons for certain actions can be explained, initially, in terms of the truth of certain reason judgement and, ultimately, in terms of the soundness of certain episodes of practical reasoning. Constructivists thus deny that reason facts are entirely independent of our mental states and activities. They are committed to the following claim:

Mind-Dependence

Facts about our mental states or activities figure ineliminably in an account of why some consideration is a reason for an agent to do something.

Note that this is a rather weak understanding of mind-dependence. It does not entail, for example, that an agent's reasons co-vary with her mental states or with how she engages in the relevant mental activities. Hence, constructivists do not suggest that practical reasons are subjective in any such straightforward sense. But they are committed to

¹⁵ See also Korsgaard (2003; 2008, 22–23).

¹⁶ That is, the practical problem schematically denoted by the concept in question; see also note 8 above.

denying that reason judgements are *objectively* true in the sense that the measure of their truth is provided by some realm of independent objects. Insofar as they want to allow for objectivity in matters of practical reasons, constructivists thus need to offer a different understanding of what that objectivity consists in, one that is compatible with their commitment to mind-dependence.¹⁷

Some authors, including both sympathisers and critics of constructivist views, have characterized them solely in terms of their commitment to the mind-dependence of the relevant facts. Russ Shafer-Landau (2003, 13), for example, introduces moral constructivism as a view which, like the realism he prefers, “endorse[s] the idea that there is a moral reality that people are trying to represent when they issue judgements about what is right and wrong”. What distinguishes constructivism, he continues, is how it conceives of this reality which moral judgements are trying to represent: it explains this reality as the “output” of a “constructive function” that takes “attitudes, actions, responses, or outlooks of persons, possibly under idealized conditions” as its inputs, whereas realists endorse the mind- or “stance-independence of moral reality” (2003, 14–15). Shafer-Landau thus portrays constructivists as *accepting* Representationalism and taking the facts represented by moral judgements to be mind-dependent in some sense. Sharon Street, an advocate of constructivism, also focuses on the issue of mind-dependence in characterizing her position, drawing upon Shafer-Landau’s notion of “stance independence” (Street 2006, 110–11).

Characterizing constructivism exclusively in terms of its commitment to mind-dependence results in a much broader understanding of the position than the one I suggested.¹⁸ In fact, the position which Shafer-Landau and those who adopt his

¹⁷ Aaron James (2012) calls this kind of objectivity “protagorean objectivity”. Carla Bagnoli (2002, 131), too, notes that constructivists are committed to an alternative understanding of objectivity that does not presuppose the accurate representation of “a (special) sector of reality.”

¹⁸ On Shafer-Landau’s characterization, dispositionalist views of value like the ones developed by David Lewis (1989) and Michael Smith (1994) as well as Mark Schroeder’s (2007) Humean reductivism about reasons seem to qualify as forms of constructivism.

characterization of constructivism seem to pick out with that label is a version of Representationalism. More specifically it is a version of Representationalism that endorses a certain conception of the kind of facts represented by our judgements: what we are “trying to represent” is a certain mind-dependent part of reality. It is thus a position that conflicts with the constructivist position presented in the present book, which rejects Representationalism altogether.

Now, as I mentioned before, everyone is free to define and use such labels as they like. The positions characterized by Shafer-Landau and myself can be understood as two alternative ways of spelling out the constructivist metaphor that reasons are not discovered but made. They emphasise different aspects of that metaphor, focussing either on the suggestion that reasons are some sort of artefact or product that is the output of a process that takes mental states or activities as an input, or on the metaphor’s implications for the relation between those reasons and our judgements about them, which is claimed not to be one of “discovery.” Nevertheless, it would be foolish to deny that both ways of characterizing constructivism are legitimate attempts to translate the constructivist metaphor into a set of more precise metaethical theses.

Still, I submit that characterizing constructivism as a position that accepts Mind-Dependence and *rejects* Representationalism carves out a more promising position that offers a genuine and fundamental alternative to the more established metaethical positions. Moreover, as we saw in the previous sections, this way of understanding constructivism takes up core elements of the views developed by Rawls and Korsgaard and is thus hardly revisionist.

Nevertheless, Rawls and Korsgaard also agree that adopting a constructivist account for a certain domain entails that the facts of that domain are mind-dependent in some sense, and that this has consequences for how we must understand the idea that truth in that domain is nevertheless an *objective* matter. In the case of Rawls, this can be illustrated with the following passage from his Dewey lectures:

Kantian constructivism holds that moral objectivity is to be understood in terms of a suitably constructed social point of view that all can accept. Apart from the

procedure of constructing the principles of justice, there are no moral facts. Whether certain facts are to be recognised as reasons of right and justice, or how much they are to count, can be ascertained only from within the constructive procedure, that is, from the undertakings of rational agents of construction when suitably represented as free and equal moral persons. (1980, 519)

This passage is often referred to when characterizing how Rawls understands constructivism,¹⁹ and it makes clear that he takes moral facts to depend on our mental states. Rawls also emphasises that constructivism and rational intuitionism come with different conceptions of objectivity. Constructivism, he stresses, is not a form of subjectivism; it allows for fallibility by providing a standard of correctness to which individual judgements can be held and which they can fail to meet. But it denies that this standard of correctness is provided by a prior and “independent order of moral values,” (1993, 112; cf. 2000, 243–47). Instead, constructivism takes a moral judgement to be correct if it “meets all the relevant criteria of reasonableness and rationality” which are incorporated into a constructivist procedure (1993, 114).

Korsgaard, too, acknowledges her view’s commitment to the mind-dependence of normative facts. As we saw earlier, she denies that a normative judgement is true because it accurately describes any normative facts that are part of some external reality. But she does not deny that there *are* normative facts, or that these facts correspond to our true normative judgements. She takes them to form a “constructed reality,” a reality constructed by the correct use of normative concepts (2003, 117). Hence, normative facts are not discovered in practical thinking, that is, thinking that is trying to solve the problems associated with those concepts. Here, Korsgaard reverses the order of explanation. According to her, normative facts can be explained in terms of

¹⁹ See e.g. Darwall et al. (1992, 138) and Street (2010, 365). Note, though, that while this passage provides a pithy statement of some important aspects of Rawls’ constructivism, it leaves others entirely unmentioned, in particular the rejection of the rational intuitionist’s representationalist account of moral judgements. Focusing solely on this passage thus bears the risk of ignoring that constructivism—understood as a comprehensive moral doctrine—is committed not only to the mind-dependence of moral facts but also, and primarily, to an anti-representationalist account of moral judgements.

practical thinking that is done correctly, that is, in accordance with the principles provided by the correct conception of practical reasons: “When those principles are applied to facts and cases, they pick out the substantive considerations that we then regard as reasons” (2008, 3). That is why, on her account, reasons are “mind-dependent”, where this means that “the body of facts in question would not exist were it not necessary for human beings to conceptualize the world in a certain way, where the aim of that mode of conceptualization is not simply one of describing the way things are” (2009a, 28).

1.5 Constructivism, explanation and the grounding relation

Constructivists, I have said repeatedly, want to *explain* reasons in terms of true reason judgements, rather than vice versa; they claim that reason facts ultimately obtain *because* or *in virtue of* the soundness of certain episodes of practical reasoning. To be able to fully understand the constructivist position and compare it to the various other accounts of practical reasons on offer, more needs to be said about what these phrases are meant to convey. What kind of explanation does the constructivist seek to provide? Reason facts are not a causal consequence of any of our judgements, so the kind of explanation at issue is not the one provided by saying that a bridge collapsed because it was hit by a truck, for example.

Instead, constructivists aim to give a metaphysical explanation of reason facts. This is the kind of explanation we might be offering when we say that the ball is red because it is scarlet, or that Peter is in pain because his c-fibres are firing, or that the vase is fragile because it is made of thin glass. We often use the phrase “in virtue of” to refer to the relation that underwrites such explanatory claims: the ball is red in virtue of being scarlet, he is in pain in virtue of his c-fibre activity, the vase is fragile in virtue of being made of thin glass. This phrase, when it is used in this way, refers to a particular ontological relation among facts: the grounding relation. Facts can be grounded in other facts, and if they are, we can give a special kind of explanation of the former in terms of the latter. The grounding relation underlies this metaphysical explanation in the same way that the relation between cause and effect underlies causal explanations. It is necessary for the corresponding explanation to be available, but it is not sufficient. As

Peter Lipton points out, “[t]he big bang is part of the causal history of every event, but explains only a few” (1990, 249). In general, whether or not a cause makes for a good explanation is highly sensitive to our interests and our background knowledge.

Similarly, the grounding relation is necessary for the availability of a distinctive non-causal kind of explanation, but whether or not giving a fact’s ground makes for a good explanation might depend on additional factors.

The grounding relation is studied and discussed extensively in contemporary metaphysics.²⁰ What motivates a large part of this attention is a shift in our conception of ontology, away from the traditional Quinean view, according to which ontology is solely concerned with what there is and aims at drawing up a large list of everything that exists, an inventory of the world, as it were. As Jonathan Schaffer puts it, Quinean ontology conceives of its object as flat, asking only whether or not it includes a particular entity, whereas his alternative, “neo-Aristotelian” conception of ontology takes reality to be structured by a relation that orders things from more to less fundamental (Schaffer 2009, 345–56). This structuring relation is the grounding relation. In addition to answering Quine’s “What is there?” ontologists thus have to answer the question “What grounds what?”

Like “existence”, the notion of grounding is a primitive metaphysical notion, one that cannot be defined in any other terms. To give an informative account of the grounding relation, we can, however, give illustrative examples, characterize its formal features, and compare it to other philosophically important relations. I have already mentioned some examples of what appear to be claims about what grounds what above. Further examples of grounding claims include “This is a triangle because it has three sides,” “Anne is married to Claire because they were pronounced spouses by a suitably authorised person,” and “There is a table in the room because there are particles arranged tablewise in the room.” In general, the relation is often taken to hold between facts concerning dispositional and categorical properties (the vase is fragile because it is

²⁰ See e.g. Fine (2001), Rosen (2010), Schaffer (2009) and the contributions in Correia/Schnieder (2012). For an overview, see Bliss/Trogon (2016).

made of thin glass), determinable and determinate properties (the ball is red because it is scarlet), macro- and microphysical objects (there is a table in the room because there are particles arranged tablewise in the room), sets and their members (there is a football team on the field because there are several players on the field), etc.²¹

In all of these cases, there is an intuitive “directedness” in the relation between the facts in question, which makes it asymmetric—if Peter is in pain because his c-fibres are firing, then it is not the case that his c-fibres are firing because he is in pain (unless we equivocate on “because” and interpret it as indicating a grounding relation on the first occurrence and an evidential relation on the second). This is true no matter which phrase we use to indicate grounding: if the ball is red in virtue of being scarlet, then it cannot be scarlet in virtue of being red, and, to use another such phrase, if being scarlet makes the ball red, then being red does not make the ball scarlet. Hence, the most salient formal features of the grounding relation is its asymmetry. Moreover, the relation is irreflexive (it makes no sense to claim that the ball is red because it is red) and transitive (if the ball is coloured because it is red, and it is red because it is scarlet, then the ball is coloured because it is scarlet).²²

How does this grounding relation compare to other philosophically important relations? One such relation is reducibility. A stated goal of many theories in all areas of philosophy, including, of course, metaethics, is to reduce certain facts to certain other facts. Is asking what grounds what just a different way of asking what is reducible to what? That depends on how exactly we understand the reductive project. Firstly,

²¹ I do not claim that the grounding relation actually holds in all of these cases, but merely that these are typically understood as instances of that relation. Of course, some of them might turn out to be false examples. As long as they help to make clear what they are supposed to be examples for, they have done their job here.

²² Many metaphysicians also introduce the notion of partial grounding: a partial ground is a proper part of a set of facts that (fully) grounds a fact, and while a fact’s full ground necessitates it, a partial ground does not (Audi 2012, 698). Here and in what follows, I am concerned with full grounding.

establishing the reducibility of, say, pain to c-fibre activity should be distinguished from providing a conceptual analysis in the narrow sense of defining “pain” in terms c-fibre activity and claiming that this conceptual relation is accessible a priori to all competent speakers. Reductivists are ultimately interested in the relations between facts or properties, not those between words or concepts (even though, of course, the latter can be a guide to the former—if “pain” can be defined as “c-fibre activity”, then pain is also reducible to c-fibre activity). The classic example of reducibility in this sense is the relation between (facts about) water and (facts about) H₂O. Note, though, that such reduction claims are often understood as establishing the identity of the properties or facts picked out by different concepts: water just is H₂O, so “this glass is filled with water” and “this glass is filled with H₂O” correspond to the same fact.

If the aim of a reductivist project is to establish an identity claim, then grounding differs significantly from reducibility, since the grounding relation is asymmetric, irreflexive and holds between distinct facts, neither of which is true of identity. I suspect, though, that very often and despite assertions to the contrary, reductivists are not actually interested in establishing identities. After all, they are usually trying to reduce some facts or properties to some “more basic” or “more fundamental” facts or properties.²³ Note also that just like “because” and “in virtue of”, “is reducible to” has a clear connotation of directedness, which suggests that what is meant by reducibility is an asymmetric relation between facts or properties that differ in their fundamentality. Grounding is just such a relation. Thus, talk of reduction is at best ambiguous and at worst misleading. To reduce A to B might simply amount to establishing a symmetric identity claim, as (perhaps) in the case of H₂O and water. But it might also turn out to mean that A-facts are grounded in B-facts, in which case we might as well say so.

²³ Sometimes, that might mean only that the same facts are better understood under another description that is in some sense more basic. In other cases, however, reduction claims are clearly intended to assert an asymmetric relation between properties or facts picked out using different concepts—this will be the case whenever someone who purports to have reduced A to B denies that B is reducible to A as well.

Another important relation that might bear some similarity to the grounding relation is supervenience. First and foremost, supervenience is a modal relation between two sets or kinds of properties: to say that the A-properties supervene on the B-properties is to say, roughly, that two things that do not differ in their B-properties cannot differ in their A-properties, either. There are several more precise specifications of this relation in the literature (see McLaughlin 1995 and McLaughlin/Bennett 2011, secs. 3.5; 4.1). But they all characterize supervenience as a reflexive²⁴ and non-symmetric²⁵ relation. Thus, the fact that certain A-properties supervene on certain B-properties does not suffice to establish the sort of directedness that is distinctive of the grounding relation, and hence does not entail that the former are grounded in the latter. But if, as is plausible, the grounding relation holds with (metaphysical) necessity,²⁶ then that relation might, conversely, offer a good explanation for why a supervenience claim holds for certain properties.

After all, supervenience claims merely describe a certain pattern among properties, they tell us nothing about why those properties exhibit that pattern. Hence, supervenience claims are usually in need of an explanation (see McLaughlin/Bennett 2011, sec. 3.7 and Horgan 1993). One possible explanation is that the two kinds of properties are actually identical. Yet, many philosophically interesting supervenience claims are made while insisting that the two kinds of properties and hence the facts about their instantiation are numerically distinct. An account of the A-properties that takes them to supervene on but be nevertheless distinct from the B-properties is often labelled a “non-reductive” account of A-properties in terms of B-properties (here “reduction” is understood as establishing identity). But it seems that establishing such a modal

²⁴ Trivially, no two things can differ in their A-properties without differing in their A-properties.

²⁵ A binary relation is non-symmetric if and only if it is neither symmetric nor asymmetric. That is, if R is a non-symmetric relation, then aRb entails neither bRa nor its negation, but is consistent with either. To see that two kinds of properties can be mutually supervenient, consider for example the radius and the surface of perfect spheres.

²⁶ For arguments to the effect that it does, see Fine (2012, 38) and Rosen (2010, 118).

connection between distinct properties does not itself explain anything or solve any problems, it rather states a phenomenon that it is the task of a genuinely informative account of the properties in question to make sense of (see Kim 1993, 165–69). The grounding relation offers a way to do so.²⁷

The preceding paragraphs, I hope, provide an understanding of the grounding relation and the sort of metaphysical explanation it affords that is sufficient for the use they will be put to in what follows. Like “existence”, “grounding” is a primitive metaphysical notion, it denotes a relation that we are familiar with from a number of contexts and examples and that we often refer to by using “because” or “in virtue of”, a relation that is asymmetric, irreflexive, transitive, and holds between distinct facts. Thus, claims about what grounds what must be distinguished from identity claims as well as claims about what supervenes on what, but they might provide an explanation for the latter. Moreover, grounding allows us to talk less ambiguously about the asymmetric, explanatorily potent dependence relation we often appear to have in mind when we think about reducibility. This notion of grounding allows us to formulate constructivism more clearly as a metaphysical view about the place of reason facts in an ontology that is structured by the grounding relation. Constructivism is a view about what grounds reason facts.

1.6 Putting constructivism on the map

Let me summarize the view as it has been outlined so far. Constructivism about practical reasons rejects a representationalist account of the nature of reason judgements. Their cognitive function is not to provide an agent with normative information by representing facts about what is a reason for doing what. Instead, they play the distinctive role of guiding the agent’s practical reasoning. Accordingly, reason judgements are not true in virtue of accurately representing any reason facts. That does not mean that they cannot be true at all. But according to the constructivist, what it is for a reason judgement to be true should rather be understood in terms of the soundness of

²⁷ See Väyrynen (2017) for a helpful overview of the difficulties that the challenge to explain supervenience raises for various other metaethical views.

the episode of reasoning that it is apt to guide. Constructivists thus deny that true reason judgements are true because certain reason facts obtain. But they do not deny that those judgements correspond to such facts, that is, they do not deny that a reason fact obtains whenever a reason judgement is true. What is crucial about constructivism, however, is that it explains the former in terms of the latter: some considerations are reasons for certain actions because the corresponding reason judgement is true, and thus, ultimately, because a particular episode of practical reasoning is sound. The constructivist's picture thus includes reason facts, but their existence is not independent of us. Instead, they are mind-dependent: facts about our mental states or activities figure ineliminably in the explanation of why they obtain.

How does this view relate to other, more established metaethical theories? It has been argued that constructivism does not constitute a *metaethical* theory at all (Darwall et al. 1992; Hussain/Shah 2006). It should be clear, however, that this scepticism is unfounded when it comes to the view presented here. It might be difficult to say exactly what conditions have to be met in order for a theory to qualify as metaethical. But I take a sufficient criterion to be that the theory answers (some of) the same questions that more paradigmatic metaethical theories (i.e., those whose credentials as metaethical are not called into question) are characteristically concerned with. Surely questions on the nature of reason judgements, on whether they can be true and what their truth consists in, and on what grounds reason facts are among those characteristically metaethical questions.

But what *kind* of metaethical theory is constructivism? In particular, is it a form of realism or anti-realism? I am not sure how useful this distinction is anymore, because it can mean so many things. If being a realist about practical reasons just means to acknowledge that we can be correct or mistaken in our judgements about them and that, if we are, these judgements are true or false in a substantive sense, then constructivists are realists. They also agree with realists that our reason judgements sometimes are true and hence that there *really are* practical reasons. Typically, however, realism is meant to be more demanding: reason facts do not only have to exist, but they have to be especially *ontologically robust*, which is often spelled out in terms of mind- or stance-

dependence (cf. Shafer-Landau 2003). A related way of drawing the distinction takes inspiration from Plato's *Euthyphro* and characterizes realism and anti-realism as disagreeing about whether there are reasons ultimately because we think there are, or vice versa. Understood in this way, constructivism is a version of anti-realism, it embraces mind-dependence and sides with Euthyphro (cf. James 2007; 2012 and Street 2010).

However, many people—particularly in contemporary metaphysics—seem to think that the question whether some particular kind of fact exists (i.e., is part of our ontological inventory) is not the only interesting, and perhaps not even the most interesting ontological question to ask. We should also be concerned with ontological structure and the grounding relation, that is, we should ask not only “What is there?” but also “What grounds what?” (Schaffer 2009). If you agree, as I do, then the classical contrast between realism and anti-realism becomes less important, and the contrast between views that do and views that do not take reason facts to be grounded in certain other facts becomes more relevant.

We can distinguish three positions on this issue. Firstly, one might take reason facts to be fundamental and not grounded in any other kind of facts; they are facts *sui generis* and make up a distinct part of reality. Furthermore, our reason judgements purport to represent those facts, and they are true, if they are true, because the relevant facts exist and they hence represent accurately. Thus, nothing grounds reason facts, and reason facts ground the truth of our reason judgements. This seems to best describe the position taken by reason fundamentalists like Derek Parfit (2011b) and Thomas Scanlon (2014).²⁸ Secondly, one might deny that reason facts are fundamental and take them to be grounded in some other kind of facts, e.g. facts about our desires, but retain the

²⁸ Note that reason fundamentalism—the claim that reason facts are neither grounded in nor identical to any other facts—must be distinguished from primitivism about the concept of a reason, which claims that this concept cannot be analysed in terms of any other concepts. Scanlon and Parfit are reason primitivists in this sense, too, but in general, those two theses can come apart. That is, one might deny the analysability of the concept of a reason while proposing a metaphysical “reduction” of reason facts (Schroeder 2007 is an example for such a position).

fundamentalist's claim about the relation between those facts and our reason judgements. That is, reason facts are themselves grounded in certain other facts, but they ground the truth of our reason judgements when they are accurately represented by them. This position is taken, for instance, by Humeans like Mark Schroeder (2007), and I will call it reason reductivism.²⁹ Finally, one might take reason facts to be grounded in our true reason judgements, so that the fact that some consideration is a reason for an agent to do something can be explained by the fact that the corresponding reason judgement is true, rather than vice versa. Of course, to avoid circularity, such a position must take the truth of our reason judgements to be grounded in something else than the obtaining of certain reason facts, and thus to consist in something else than their accurate representation. This is the constructivist view as I introduced it.

Note that Humean and other versions of reductivism that take reason facts to be grounded in psychological facts, for instance about the agent's desires, also entail that reasons depend in some way on our mental states or activities. Focussing solely on a position's stance on mind-dependence will thus lead you to assimilate constructivism to reductivism and to ignore the important difference that distinguishes constructivism from both fundamentalism and reductivism: constructivists reject the idea that the truth of a reason judgement is grounded in the existence of a fact that it represents as obtaining. This is why Representationalism is at the heart of the issue between constructivism and its alternatives.

This brings us to a worry about constructivism that has been raised repeatedly in the literature: is constructivism really a *distinct* metaethical view, or does it, upon closer inspection, collapse into one of the more familiar views?³⁰ Those who think of

²⁹ I choose this term despite the ambiguity noted in the previous section because it is so widely established in the debate. As I said above, I suspect that many of those who are considered—or consider themselves—reductivists about reasons actually intend to argue for a grounding claim. For the sake of completeness, however, I will classify anyone who argues for the strict identity between facts about reasons and certain other facts as a reductivist as well.

³⁰ Versions of this worry are raised, e.g., by David Enoch (2009), James Lenman (2012), and Michael Ridge (2012).

constructivism primarily as a view that rejects Representationalism often question whether it offers a genuine alternative to expressivism. As I have outlined in section 1.2, however, while constructivism and expressivism share a commitment to rejecting Representationalism, they differ substantially in how they conceive of the role that reason judgements play in an agent's psychology and, consequently, in their positive accounts of the nature of those judgements. Those who think of constructivism primarily as a view that embraces mind-dependence, on the other hand, tend to question whether it offers a genuine alternative to familiar reductive accounts. David Enoch (2009, 328–29), for example, notes that dispositionalist views like David Lewis' (1989) qualify as a version of constructivism as he characterizes it. But as we just saw, if we understand constructivism along the lines suggested here, it differs significantly from those reductive accounts, because it reverses the order of explanation between the truth of reason judgements and the obtaining of the facts to which they correspond.

In addition to this “someone already said that”-version of the distinctness worry, Enoch also presents a more general version of the worry:

It seems to be possible to divide up the logical space of metaethical positions rather neatly, using a series of yes-no questions. [...] In such a way, we seem to get an exhaustive classification of all logically possible metaethical positions. But when the details of this little exercise are filled in [...], no room remains for a distinctively constructivist position. (2009, 329)

His idea is that by asking constructivists to answer a series of yes-or-no questions, one can usher them along the forking paths of a taxonomic tree, all of which eventually lead to boxes that already have different and familiar names on them (such as “Dispositionalism”). We can call this the argument from taxonomic exhaustion. Does it establish that constructivism is not a distinct metaethical view?

It does not. Firstly, the most we can and should say about metaethical views that end up in the same box of such a taxonomy is that their disagreement on one or more of the many issues discussed in metaethics is sufficiently small to allow them to answer certain questions in the same forcedly non-nuanced way. Whether or not this should prompt us to classify them as versions of the *same* position rather than as distinct

alternatives, it seems to me, is not a particularly sensible question to ask, as long as there is some substantive element of disagreement. Moreover, as we saw in the case of “realism” above, such classifications are always prone to lead to merely verbal disputes. They derive the little importance they have from the much more important goal of finding out which answers to the underlying questions are correct, and their adequacy is highly sensitive to the context and the background assumptions of the discussions in which they are used. This is not to deny that taxonomies can be very useful, say, to give an overview of an area or to make certain dialectical situations explicit. But we should keep in mind that what questions we use in generating them will depend not only on abstract systematic considerations, but on what questions we deem important enough to give different names to different answers, which in turn depends on what problems we deem worth solving and what general concerns we take metaethical theories to be in the business of addressing.

Secondly, even if we put these general concerns aside, Enoch’s argument still fails to establish that there is “no room [...] for a distinctively constructivist position” in metaethics (2009, 329). This is because the taxonomy he offers is not exhaustive, it conflates important distinctions in the metaethical realm. In particular, it moves from answering “yes” to the question “Are any moral sentences strictly speaking true?” directly to the question “Are moral properties (or facts) naturalistically reducible?” (ibid.) Thereby, it ignores precisely those questions concerning the proper understanding of truth and the relation between facts and true judgements on which constructivists and their opponents disagree. Moreover, it does not include any questions on the nature of the normative judgements in question. Rather than demonstrating that there is no room for constructivism in metaethics, Enoch’s argument thus illustrates how a one-sided conception of the issues that metaethicists should be concerned with can make us blind to the very possibility of a distinctively constructivist account of practical reasons.

1.7 Plan of the book

In the remainder of this book, I will develop the constructivist account of practical reasons that I outlined in the preceding sections into a fully-fledged metaethical view,

make its theoretical commitments explicit and respond to the most pressing worries to which it gives rise. My aim in doing so is not to convince you that constructivism is the best metaethical view, or that it is superior to any of its alternatives. In fact, I will have little to say about those alternatives, except where I need to distinguish constructivism from them. My aim is a more modest one: I want to show that constructivism offers a distinct, coherent and *prima facie* attractive account of practical reasons, and to make as clear as possible what accepting that account would involve. This, I hope, will facilitate a more focussed and in-depth debate about its merits.

To that end, the following chapters will elaborate the aspects of the constructivist view that I outlined in sections 1.2 through 1.4 and address the two worries that have already been mentioned—concerning the view’s distinctness and circularity—in much greater detail.

Chapters 2 and 3 will develop the constructivist’s alternative to Representationalism. Chapter 2 argues that reason judgements play an important and essential role in practical reasoning: they guide the reasoning process. Chapter 3 then shows how this role can be used to formulate an anti-representationalist account of the nature of reason judgements. It also discusses how that account differs from the expressivists’ understanding of the nature of such judgements, who also reject Representationalism.

Chapter 4 presents an account of the truth of reason judgements that does not presuppose the prior existence of reason facts but nevertheless understands truth to be a substantive property that can bear the explanatory weight that the constructivist account of practical reasons puts on it. In particular, it suggests that constructivists can formulate such an account by adopting a theory of truth that has been developed most prominently by Crispin Wright (1992) and Michael Lynch (2009).

According to the account of truth presented in chapter 4, the truth of a reason judgement is ultimately grounded in the soundness of the episode of practical reasoning that it is apt to guide. Most importantly, such an episode is sound only if it is an instance of correct reasoning. That is why chapter 5 addresses the question of what it is for practical

reasoning to be correct. To avoid circularity, constructivists must answer this question without presupposing any truths about practical reasons. I show that this can be accomplished by taking the standard of correctness for practical reasoning to be provided by the constitutive rules of that activity. Like the rules of chess, these rules necessarily govern the activity even though there need not be any reason to comply with them.

Finally, chapter 6 distinguishes the particular version of Mind-Dependence that constructivists endorse from the version to which reductivists like Mark Schroeder (2007) are committed. These versions differ in the resources they offer for accommodating the idea that practical reasons are objective in some sense. I thus discuss how constructivism and its alternatives fare in accommodating the fallibility of our reason judgements, the fact that some reasons are universally shared, as well as the modal robustness of our reasons, which are three rather different issues associated with the ambiguous notion of objectivity.