

Towards a Semantics for Metanormative Constructivism

Abstract: The status of constructivism as a metaethical or metanormative theory is unclear partly due to the lack of a clear semantics for central normative terms such as ‘reason’ and ‘ought’. In a series of recent papers, Sharon Street has attempted to clarify the central commitments of constructivism by focusing on the idea of a practical point of view and what follows from it. We improve upon the informal understanding provided by Street and attempt to provide a semantics for ‘ought’. Our semantics respects the core intuition of the constructivist that normative claims are made true because of our practical commitments as agents and also reflects the constructivist’s commitment to the centrality of practical deliberation to normative truth. On our view, a normative claim of the form ‘ Γ S ought to ϕ \neg ’ is true if ϕ is entailed from S’s set of evaluative attitudes. We argue that a virtue of our definition is that it allows us to see precisely what is distinctive about constructivism as opposed to realism and expressivism.

The core intuition of the constructivist is that normative claims are made true because of our practical commitments as agents. Often this is contrasted with a certain kind of realism, which holds that normative claims are true independently of any of our commitments (Wallace 2012: 23, James 61). But it has also been argued that constructivism is itself a form of realism (Hussain and Shaw 2006: 270, Copp 2013). Alternatively, constructivism has been seen as a form of expressivism (Gibbard 1999, Lenman 2012, Ridge 2012: 157) and defenders such as Korsgaard have even argued it is compatible with both (Korsgaard 2003: 122). Since constructivists often neglect traditional questions about moral language, arguably, it isn’t even a metaethical theory at all but rather, can be better understood as a substantive normative ethical theory (Hussain and Shaw 2006: 270, Wallace 2012: 23).

It is plausible that the primary reason that constructivism’s status as a theory is so unclear is because of its lack of a clear semantics for normative terms such as ‘reason’, ‘ought’, ‘right’, and ‘wrong’. Once the constructivist is pressed to answer some traditional metaethical questions,

the view might well collapse into some other well-known view. This is what leads Michael Ridge to complain that constructivists have failed to provide “an approach to metaethics which is at once thoroughgoing, novel, and plausible” (Ridge 2012: 138). In this paper, we aim to do just that.

We will take as our point of departure, Street’s well known “practical standpoint characterization” of constructivism which holds that the truth of a normative claim consists in that claim’s following from the practical point of view (Street 2010: 364). Although we will follow Street in emphasizing what follows from the point of view of agents deliberating about practical problems, we will show that as stated, there is a crucial confusion about the kind of claim that follows from the practical point of view. At first glance, the confusion will seem small and easy to fix, but we will show that resolving the confusion suggests a complete re-understanding of the relationship between normative claims and the practical point of view. In the second section, we will articulate and defend the suggested re-understanding. According to our view, practical claims must be distinguished from normative claims, but they are related to each other in much the same way that premises of an argument are related to claims about what follows from these premises. In the final section, we defend this as an interpretation of constructivism by showing its compatibility with the core intuition as well as its plausibility and novelty. Since we focus here on providing a semantics for ‘reason’ and ‘ought’ and leave questions of the proper understanding of ‘right’ and ‘wrong’ for another day, we call our view a metanormative as opposed to a metaethical constructivism.

Street’s standpoint constructivism

While it is unclear if Street intends her account as a comprehensive semantics for constructivism, over the course of a series of papers (Street 2008, 2010, 2012) she attempts to give an account of

sentences of the form $\ulcorner S \text{ has reason to } \phi \urcorner$ and indirectly, $\ulcorner S \text{ ought to } \phi \urcorner$. Street (2010) develops her account by asking us to consider the example of Ann who values counting blades of grass above all else but does not see any reason to buy a calculator. According to Street, the normative claim “Ann has reason to buy a calculator” follows from within Ann’s evaluative point of view. Generalizing from this case allows Street to provide the following truth conditions for normative claims:

According to constructivist views in ethics, the truth of a normative claim consists in that claim’s being entailed from within the practical point of view. (Street 2010: 367)

Presumably, then since Street thinks that “Ann has reason to buy a calculator” is a true normative claim, she must think that the claim follows from within the practical point of view. However, the two most natural understandings of “the practical point of view” in this context are Ann’s own practical point of view and the practical point of view of a generic agent, and as we will show, “Ann has reason to buy a calculator” and other claims about what Ann has reason to do simply do not follow on either of these understandings.

According to classic statements of constructivism, like Rawls’ and Korsgaard’s, the practical point of view differs from the theoretical point of view because it seeks not to describe the world but to offer “solutions to practical problems” (Korsgaard 2003: 115, Rawls 1999). On this understanding, when we deliberate about what to do, we enter into a mode of deliberation that is quite different from when we deliberate about what to think. The result of such reasoning is not a revision in our beliefs but in our intentions or plans. Whereas theoretical reasoning is guided by epistemic norms, practical reasoning is guided by practical norms like means-ends inference rules. A solution to a practical

problem *follows* from an agent's practical point of view, presumably if it is entailed by the intentions and commitments of an agent via the norms of practical reasoning. Unfortunately, however this natural understanding of what follows from the practical point of view cannot be what Street has in mind because on this interpretation normative claims like “Ann has reason to buy a calculator” cannot follow from the practical point of view. Rather, the conclusion must be something that Ann could intend to do, like *to buy a calculator*.¹ Although a seemingly small point, this will be crucial for our overall analysis, so it is worth going slow here.

Imagine that Ann learns that calculators are effective tools for summing the counts from previous days and then she puts this fact together with her commitment to counting blades of grass. She then practically deliberates and decides to buy a calculator. When Ann is deliberating she is deliberating about blades of grass and their relationship to calculators; she is not deliberating about herself or about what reasons she has.

Compare these two practical arguments:

Deliberation to a practical conclusion

P1. Buying a calculator would help with counting blades of grass

Deliberation to a normative conclusion

P1. Buying a calculator would help with counting blades of grass

¹ Here we follow Gibbard 2003: 47 in thinking that the contents of a belief are different than the contents of a decision (or a plan) and marking this difference by giving the contents of decisions the infinitive form. As Gibbard notes, in doing this, we thereby reject the force-content distinction, which holds that plans and beliefs have the same content but differ in force.

P2. To count blades of grass

P2. To count blades of grass

Conc: To buy a calculator

Conc: Ann has reason to buy a calculator

The premises are the same in each case - but only the argument on the left is a valid practical argument. In the argument on the right, the conclusion does not follow from the premises and so the conclusion does not follow from within Ann's practical point of view. The conclusion of her deliberation is not a claim about what she has reason to do. Instead, the conclusion of her deliberation is the thing that she, as a result of her deliberation, now has reason to do. We conclude that claims of the form $\ulcorner S \text{ has reason to } \phi \urcorner$, $\ulcorner S \text{ ought to } \phi \urcorner$, and other related forms are typically not the conclusions of S's own inferences and therefore cannot follow from within S's point of view.

Although Street suggests that the normative claim is supposed to follow from within Ann's own practical point of view, her actual definition of normative truth talks about "*the* practical point of view". Perhaps, then we are trying to derive the claim, "Ann has reason to buy a calculator", from the wrong thing. *The* practical point of view might mean the point of view of a generic rational agent, thought of as an ordinary rational agent who has had all of her contingent preferences stripped away. On this reading, the normative truths are the normative claims that can be the conclusion of a practical deliberation for any agent at all. But this option is clearly less plausible than the one that started from Ann's evaluative attitudes. If the claim, "Ann has reason to buy a calculator" doesn't follow from Ann's evaluative attitudes, it certainly isn't

going to follow from a generic agent's attitudes who is thought to have fewer preferences even than Ann.

There are indications that Street doesn't think of the practical point of view in either of the natural ways that we have articulated. After discussing various *restricted* forms of constructivism, Street says:

According to **thoroughgoing** or **metaethical constructivist views**, the truth of a normative claim consists in that claim's being entailed from within the practical point of view, where the practical point of view is given a formal characterization (Street 2010: 369).

Here, the notion of a "formal characterization" plays a special role:

"To give the practical point of view a formal characterization is to give an account of the standpoint of valuing or normative judgment as such, where this involves giving an account of the attitude of valuing that does not itself presuppose any substantive values but rather merely explicates what is involved in valuing anything at all." (369)

When Street talks about what follows from the nature of valuing as such, perhaps she has in mind high-level conceptual analysis of the idea of valuing. So just as "aiming at the truth" might follow from a high-level conceptual analysis of the nature of belief as such, perhaps there are similar high-level claims that follow by means of conceptual/philosophical analysis of the nature of valuing as such. Let us call these high-level features "structural features" of valuing. What follows from these structural features is unclear. But for now, it is enough to note that it is clear that nothing about Ann in particular could follow from these structural features.

Street recognizes the problem in a long footnote in which she considers the objection that her view doesn't seem to be able to show the truth of individual normative claims, like "Ben has

normative reason to add to his stamp collection.”² In the footnote, Street concedes that normative claims about Ben do not *directly* follow from the practical point of view as such. But she then claims what does follow directly are structural claims like “a valuing creature’s normative reasons depend in a certain way on that creature’s contingent evaluative attitudes” or conditional claims like “*given a certain set of contingent evaluative attitudes* – one that probably includes, among other things, a love of stamp collecting – an agent has normative reason to add to his stamp collection” (Street 2010: 282). It is not obvious how this is supposed to help, but perhaps the idea is that specific claims like “Ben has normative reason to add to his stamp collection” do follow when we add various empirical assumptions (such as that Ben loves stamp collecting). Starting with the practical point of view of agency as such, we might argue to a structural claim about the dependence of reasons on evaluative attitudes. From that we might get a conditional claim like *given a love of stamp collecting, an agent has normative reason to add to his stamp collection*, and then by adding the empirical claim about Ben’s love of stamp collecting, we derive, finally, that Ben has normative reason to add to his stamp collection.

But even if that works, it should be clear now that we are dealing with two very different understandings of ‘the practical point of view’ and ‘follows’. According to what we called the “natural understanding,” the practical point of view starts with Ann’s or Ben’s intentions and commitments and arrives at other intentions and commitments by means of practical reasoning. On such an understanding of the practical point of view, the sorts of claims that would follow from Ann’s or Ben’s practical point of view would be claims like “to buy a calculator” or “to add to my stamp collection” since these are the sorts of things that Ann or Ben could intend. Since

² Notice here that Street has added the word ‘normative’ into the normative claim itself. It is doubtful that this makes a difference in the overall argument. For example, it is just as implausible that claims about what normative reasons Ben has follows from the practical point of view as is it that claims about what reasons Ben has follows.

normative claims like “Ann has reason to buy a calculator” or “Ben has normative reason to add to his stamp collection” are not the sorts of thing that Ann or Ben intends, they simply would not follow from the practical point of view.

If on the other hand, we follow what seems to be Street’s derivation of a normative claim in the Ben case, then we arrive at a very different picture. *The* practical point of view involves high-level theoretical claims *about* the nature of valuing as such. Ann’s or Ben’s intentions and commitments do play a role but only as empirical premises *about* what Ann or Ben intend or value. The normative claims themselves follow from a combination of the high-level theoretical claims about the nature of valuing as such combined with the empirical claims about what Ann or Ben values. Moreover, these normative claims do not follow by means of practical norms like instrumental reasoning, but instead directly follow from the ordinary rules of logic. In fact, on this understanding of ‘follows’ and ‘practical point of view’, the only thing practical about this reasoning is its subject matter. The reasoning itself, in its premises, conclusions, and the norms that guide it is all theoretical.

Street herself seems to be unsure of which of these two pictures she is endorsing. On the one hand, her conviction that claims like “Ann has reason to buy a calculator” could be derived from the practical point of view suggests she must have the theoretical story in mind, but other times when she talks about what is entailed from “within” Ann’s practical standpoint, it seems like she must have the practical story in mind. In either case, Street seems to be stuck with a dilemma: if she accepts the practical story she stays close to the original constructivist emphasis on practical reasoning but normative claims like “Ann has reason to buy a calculator” do not follow from within Ann’s practical point of view. If, on the other hand, she accepts the theoretical story, then perhaps she can get the normative claim to follow, provided that she adds

a whole bunch of empirical claims about what Ann intends, but then normative truth is not really a matter of practical reasoning at all, but rather theoretical reasoning about our practical commitments. But this picture strays too far from classic statements of constructivism, which view normativity as stemming directly from practical deliberation. Therefore, in what follows, we pursue the first alternative—making the idea of what practically follows from Ann’s own commitments central, but giving up on the idea that normative claims themselves are what follows.

A semantics for a truly practical constructivism

We have seen that claims of the form $\ulcorner S \text{ has reason to } \phi \urcorner$ do not follow from within an agent’s practical point of view. But surely something related to this follows. If we examine a normative truth of the form $\ulcorner S \text{ has reason to } \phi \urcorner$, we should be asking whether the practical judgment $\ulcorner \text{to } \phi \urcorner$ follows and not whether $\ulcorner S \text{ has reason to } \phi \urcorner$ itself follows from anything. This suggests a (seemingly) small revision to Street’s definition of normative truth.

Normative Truth (Street): The truth of a normative claim, $\ulcorner S \text{ has reason to } \phi \urcorner$ consists of this claim being entailed from within the practical point of view.

Normative Truth (Revised): The truth of a normative claim, $\ulcorner S \text{ has reason to } \phi \urcorner$ consists of $\ulcorner \text{to } \phi \urcorner$ being entailed from within the agent S’s practical point of view.

Although we think that this revision is on the right track, the ‘entailed’ is too strong for sentences of the form $\ulcorner S \text{ has reason to } \phi \urcorner$. For example, if I wanted to get to Amarillo and Amarillo is (far) north of where I currently am, I may have a reason to start walking north, but walking north

is not *entailed* from my practical commitments. After all, I could (more sensibly) drive there. This suggests that if we want to use ‘entails’ and ‘conclusion of a deliberation’ as part of our explanans, we need to change our explanandum. Instead of capturing the meaning of sentences of the form ‘ Γ S has reason to ϕ ’, our revised definition comes closer to capturing sentences of the form ‘ Γ S ought to ϕ ’. In fact, we say that Ann *ought* to buy a calculator if her attitudinal set, taken as a whole, positively settles the question of whether to buy a calculator. In the simple case, this may be because buying a calculator is necessary for Ann to achieve one of her ends, or in more complex cases, because buying a calculator is the uniquely best solution to a practical problem faced by Ann. This suggests the following definition:

Normative Truth (ought) (NTO): The truth of a normative claim, ‘ Γ S ought to ϕ ’ consists in ‘ Γ to ϕ ’ being practically entailed from the elements of S’s attitudinal set.

Consider, once more, the case of Ann. Let us say that her attitudinal set is composed of a mixture of plans and beliefs, and let us distinguish the two by writing plans as infinitive verb phrases (of the form ‘ Γ to ϕ ’ and beliefs as descriptive sentences.³ In this case, Ann’s attitudinal set will have the following elements: $A = \{to\ count\ blades\ of\ grass,\ buying\ a\ calculator\ is\ necessary\ for\ my\ counting\ blades\ of\ grass,\ buying\ a\ calculator\ is\ within\ my\ power^4,\dots\}$ According to NTO, the

³ As will shortly become clear, our semantics are heavily influenced by Gibbard’s own fact-prac worlds (Gibbard 2003). In a similar vein, Peter Vranas has argued in a series of papers that there is a logic of imperatives despite the fact that imperatives are neither true nor false (Vranas 2008, 2010, 2011, 2016).

⁴ As many other have noted, going back at least to Aristotle, means-ends reasoning requires that the means are within our power. In Broome’s phrase, it requires that the issue is “up to” us (Broome 2015: 162).

normative claim *Ann ought to buy a calculator* is true just in case *To buy a calculator* is practically entailed by elements in *A*.

Giving a complete theory of practical entailment is a difficult task. Here, we give only a rough sketch of what such a theory might look like. One natural way is to give a rules-based syntactical understanding of entailment. If Γ is a set of propositions, then Γ syntactically entails Γ to ϕ (\vdash_p) just in case it is possible to reach Γ to ϕ starting from elements of Γ and proceeding only by the repeated application of valid practical rules such as a Means-Ends Inference Rule. For example, in Ann's case, we say Ann's set $A \vdash_p$ *to buy a calculator* because *to buy a calculator* follows by the Means-Ends Inference Rule from the elements of Ann's set: $\{to\ count\ blades\ of\ grass,\ buying\ a\ calculator\ is\ necessary\ for\ my\ counting\ blades\ of\ grass\}$. Under this interpretation, NTO would claim that if it is in fact the case that $A \vdash_p$ *to buy a calculator*, then the normative proposition *Ann ought to buy a calculator* is true.

Another natural way to formulate practical entailment is to model it after semantic entailment. First, we say that an agent's intention Γ to ϕ is satisfied in a world if and only if the agent performs that action in that world. We can now define practical entailment in the natural way: Γ practically entails (\models_p) Γ to ϕ just in case in every world in which all the descriptive elements of Γ are true and all the practical elements of Γ are satisfied, then Γ to ϕ is also satisfied. So, for example, *To buy a calculator* is entailed by Ann's set A because in all worlds where it is true that 1) Ann's buying a calculator is necessary for Ann counting blades of grass,⁵ 2) it is true that buying a calculator is up to Ann, and 3) Ann's intention to count blades of grass

⁵ In order for "Ann's buying a calculator is necessary for Ann counting blades of grass" to be true in a world, we need only that this claim is as least as strong as the material conditional for our entailment claims to hold.

is satisfied, then her intention to buy a calculator is also satisfied. NTO claims that if such an entailment relation holds between the elements of Γ and $\ulcorner \phi \urcorner$, then it is true that $\ulcorner S$ ought to $\phi \urcorner$.

Of course, in ordinary situations, it is rarely the case that a unique course of action is deductively entailed by our set, and since we often still want to say that there are things that we ought to do in these cases, it is useful to relax our understanding of the phrase ‘S ought to ϕ ’. Often, this phrase is not about deductive entailment but rather a looser sense of entailment where $\ulcorner \phi \urcorner$ is entailed because it is the *best* of the available alternatives. While walking, hitchhiking, and taking the bus might all be ways of getting to Amarillo, given other desires for convenience and safety, driving might be the best way. If we become convinced of the superiority of driving, then we might well say that we ought to drive to Amarillo and ought not to walk or take the bus even though each of these other ways might indeed get us, albeit imperfectly, to Amarillo. This suggests that it might be useful to talk about practically optimific options where an option is practically optimific if it maximally satisfies the agent’s values and beliefs. This might be formalized by invoking a ranking function, Δ , that takes as its inputs an agent’s deliberative set and a list of options and has as its output a ranking of these options. Γ would optimifically entail $\ulcorner \phi \urcorner$ if and only if ϕ -ing outranks all other options according to Δ . In other words, if there were some ranking function Δ which ranks options according to (for example) how many of the agent’s (strong) values are satisfied, then we could say that an option, P, is optimific relative to a set of values if for any alternative option Q, P is ranked higher than Q – that is, $\Delta(Q, \Gamma) < \Delta(P, \Gamma)$. We could then define $\models_{p, \text{optimific}}$ as follows: $\Gamma \models_{p, \text{optimific}} \ulcorner \phi \urcorner$ just in case $\ulcorner \phi \urcorner$ is contained in every optimific option, i.e. one’s ϕ -ing is satisfied in every optimific outcome according to Δ . Since driving north may very well be included in every optimific solution to S’s

practical dilemma, we would say that S's evaluative set $\Gamma_s \models_{p, \text{optimific}} \text{to drive north}$. On the other hand, it is not true that $\Gamma_s \models_{p, \text{optimific}} \text{to walk north}$ precisely because there are some optimific alternatives that do not involve walking north. That the practically optimific entailment relation holds in the case of driving north means that the sentence $\ulcorner S \text{ ought to drive north} \urcorner$ is true. That the entailment relation fails in the second case means that it is not true that $\ulcorner S \text{ ought to walk north} \urcorner$. If we understand driving and walking here to be incompatible, then it will also be true that $\ulcorner S \text{ ought not to walk north} \urcorner$. Note that in the case of ties, then neither of the options will be practically required, though their disjunction will be.

Of course, our ranking function will probably not be specified very precisely. All that we insist on here is that some rough ranking function is implicit in deliberation and that this ranking function is independent of the set of values that it ranks. We state this second clause explicitly because it is important to understand the sense in which 'ought' claims and the entailment relations they are about do and do not depend upon the values of the agents. Of course there is one sense in which what is entailed by the agent's evaluative set depends on what is in that very set – this is the same sense in which the conclusion of any argument depends on its premises. But importantly, *that* the conclusion follows from a given set is not dependent on that very set. In the same way, whether a given agent's evaluative set entails $\ulcorner \text{to } \phi \urcorner$ is completely independent of whether that agent thinks that it does.

Even though the semantics provided is only a rough sketch, it should be obvious that we have modeled our account on the theoretical case. On our semantics, $\ulcorner S \text{ ought to } \phi \urcorner$ means that $\ulcorner \Gamma_s \models_p \text{to } \phi \urcorner$. Because of the parallel between ' \models_p ' and ' \models ', we should expect that sentences of

the form $\ulcorner S \text{ ought to } \phi \urcorner$ share a number of properties with sentences of the form $\ulcorner \Gamma \text{ entails } B \urcorner$. Here are three such properties:

First, note that claims about entailment are not (typically) the premises or conclusions of arguments. When we say that $\ulcorner \Gamma \text{ entails } B \urcorner$ is true, if $\Gamma = \{A, A \rightarrow B\}$, then A and $A \rightarrow B$ are premises and B is the conclusion. The word ‘entails’ does not occur in the premises or the conclusion. Similarly, when $\ulcorner S \text{ ought to } \phi \urcorner$ is true, this represents the validity of a certain practical argument. We should not expect normative claims themselves to appear as premises or as the conclusion of such arguments.

Second, $\ulcorner \Gamma \text{ entails } B \urcorner$ is true if a certain relationship holds between B and the members of Γ . Similarly, $\ulcorner S \text{ ought to } \phi \urcorner$ is true when a certain relationship holds between $\ulcorner \text{to } \phi \urcorner$ and the elements in the set of S ’s attitudes. Since the elements in S ’s attitudinal set are first-order judgments such as *to count blades of grass* and *buying a calculator is necessary for my counting blades of grass*, normative claims are better understood as second-order judgments which are made true in virtue of the relationship between first-order judgments. Since practical judgments are first-order judgments, we can now see that normative judgments are not a species of practical judgment.

Third, when it is true that $\ulcorner \Gamma \text{ entails } B \urcorner$, no property is predicated of B . Similarly, we should not expect that any normative property—such as the property *ought to be done*—is predicated of ϕ just because $\ulcorner S \text{ ought to } \phi \urcorner$ is true.

Street herself wants to explain sentences of the form $\ulcorner S \text{ has reason to } \phi \urcorner$. We noted above that these sentences are considerably weaker than sentences of the form $\ulcorner S \text{ ought to } \phi \urcorner$. Taken in the broadest sense, $\ulcorner S \text{ has reason to } \phi \urcorner$ if there is anything at all in her attitudinal set

that could explain her ϕ -ing. One obvious and popular analysis of this claim is to rely on an analogy between reason and a weak sense of evidence, for example, one in which X is evidence for Y if X raises the probability of Y – i.e., $P(Y|X) > P(Y)$. This suggests the following definition:

Normative Truth (reasons): The truth of a normative claim of the form ‘ S has reason to ϕ ’ consists in there being a state of affairs P and an action ϕ , such that P is something S values from within her practical point of view and ϕ is such that ‘to ϕ ’s being satisfied makes P more probable.

Of course there are hard questions about the semantics of probability claims like these, but these questions can be safely ignored here.⁶ All we really need is that like the case of ought, the truth of ‘ S has reason to ϕ ’ is independent of whether the agent thinks they have such a reason.

In the original Ann case, P would be the state of affairs that she aims at when she aims to count the grass and ϕ would be buying a calculator. P is valued from within her practical point of view in the sense that it is something that she intends to bring about. Since buying a calculator makes it more likely that she succeeds in counting the grass, she has reason to buy a calculator. Similarly, since walking north makes it (slightly) more likely that I get to Amarillo, I have reason to walk there even if it is far better, all things considered, to drive. In this sense, P is a reason for S to ϕ , if P can explain why, according to S ’s lights, S ϕ s.

⁶ Defenders of probabilistic accounts of desire promotion as an account of reasons include Finlay (2006) and (2014), and Schroeder (2007), but see Behrens and DiPaulo (2011) and Sharadin (2015) for worries about these probabilistic accounts.

Although the above semantics for $\ulcorner S \text{ ought to } \phi \urcorner$ and $\ulcorner S \text{ has reason to } \phi \urcorner$ are just a sketch, they are sufficient, we believe, to draw the following important lesson. On our understanding of constructivism, it is crucial to distinguish between a *practical judgment* and a *normative judgment*. A practical judgment has the form $\ulcorner \text{to } \phi \urcorner$ and a normative judgment has the form $\ulcorner S \text{ ought to } \phi \urcorner$. The former are the elements of practical reasoning, have world-mind direction of fit, are first-order commitments, and are things that can be practically endorsed by a subject. The latter are statements about practical reasoning, have mind-world direction of fit, are second-order commitments, and the sense in which they are endorsed by a subject must be different from a first-order commitment lest they reduce to such a commitment.

Thoroughgoing, Novel, and Plausible

We have now given an account of the truth of normative claims such as $\ulcorner S \text{ ought to } \phi \urcorner$ where the truth of such claims is grounded in the evaluative commitments of S . Our task was to provide a constructivist semantics of such judgments. Have we succeeded? Recalling Ridge's complaint about constructivists, our account is novel in that we know of no other account like it and is plausible in so far as it explains basic cases such as Ann's. However, we have said nothing so far about Ridge's third constraint, thoroughgoingness.

A restricted constructivist holds that the truth of one domain of normativity, call it the target domain, can be constructed from another domain of normativity, call it the source domain. Thus, Street views Rawls as constructing truths of political justice from normative truths about "fair bargaining conditions and the freedom and equality of persons" (Street 2010: 368). Similarly, Street views Scanlon (1998) as constructing truths about right and wrong from truths about reasons—such as the truth that we have "reason to live with others on terms that no one

could reasonably reject” (Street 2010: 368). As Street herself notes (368), it is clear from the way restricted constructivism works that one can be a constructivist about the target domain and be a realist or expressivist about the source domain. In other words, one could be a metapolitical constructivist (a constructivist about justice) or metaethical constructivist (a constructivist about right and wrong) without being a metanormative constructivist (a constructivist about reasons generally). Indeed, Scanlon himself is a constructivist about moral truths but a realist about the truths of normativity from which the moral domain is constructed.

On the face of it, our account seems to be thoroughgoing since we have aimed to give an account of all normative sentences of the form “ S ought to ϕ ” and “ S has reason to ϕ ” and not just a particular kind of normative sentence like “ S has political reason to ϕ ” or “ S ought morally to ϕ ”. But perhaps this thoroughgoingness is illusory. To see why an objector might think this, recall our overall account. We start with the evaluative attitudes of an agent. Practical judgments which follow from this set of attitudes are then said to be normatively required. But some might worry that we smuggle in an additional normative feature when we talk about practical entailment. The worry can be put most forcefully if we consider a case of actual deliberation like that of Ann’s. What allows for Ann to conclude to the intention of buying a calculator is that this action is a means to her end. But, asks the objector, what about the means-ends principle itself? Surely, this principle is itself normative. As Wallace puts it, “the question that immediately suggests itself about the role of rational procedures in the constructivist account is: What accounts for their authority over the agent?” (Wallace 2012: 33). The constructivist seems to be left with three bad options: Either 1) the principles of deliberation are themselves constructed from the agent’s values 2) the principles of deliberation are just unexplained normative truths or 3) the principles of deliberation are not themselves normative. 1) is the

alternative Wallace himself thinks that constructivists ought to go with, but he then rightly points out that what follows from what is then inappropriately contingent. 2) seems to turn the account into a form of local constructivism where the source domain is the principles of deliberation and the target domain is other normative sentences. Unfortunately, the source domain will turn out to consist of unconstructed normative claims that have a realist status. Finally, 3) seems to invite the objection that these principles of deliberation will not be authoritative for us. If they are not normative, then why should we obey them in the first place?

If we are right, then Wallace is too quick to reject 3). In our version of constructivism, normative claims are all of the form $\ulcorner S \text{ ought to } \phi \urcorner$ or $\ulcorner S \text{ has reason to } \phi \urcorner$. These normative facts are made true because of practical entailment relations that are not themselves normative. These are not normative because our evaluative set practically entails an action ϕ , if and only if a certain non-normative fact obtains, for example, in all worlds in which our beliefs are true and our intentions are satisfied, $\ulcorner \text{to } \phi \urcorner$ is also satisfied. Whether such a relationship holds between worlds is not a normative question at all. The same basic point holds whether we understand practical entailment as optimizing or as a syntactic relation. So we have non-normative facts about practical entailment, and we have the normative facts that apply to particular agents in virtue of them. Nothing else is needed.

A critic might insist that something else is needed. Perhaps some normative principles, normative requirements, or normative standards of correctness. Although the burden is on the critic to explain what is missing from our account, here we will make three such attempts on behalf of the critic and show that none of them succeed.

First, if the practical entailment facts are not themselves normative, what if we simply ignore them? What if we fail to take the means to our ends? In such a case, we have violated a

norm. But the norm is not “You must take the means to your ends” or anything like it. Rather, if you have an end such as *to get to Amarillo* and you have a belief such as *going north is necessary for me to get to Amarillo*, then as a matter of fact, your evaluative attitudes entail *to go north*. Thus it is true that you ought to go north. So if you don’t take these means and don’t go north, you have violated a norm that applies to you. This requires no additional normative principles.

A second try might focus on the means-ends principle directly and ask what its status is. But what exactly does such a principle say? To be sure, there are lots of slightly distinct variations in the neighborhood. One plausible version is “If you intend to ϕ and believe that Ψ -ing is necessary for ϕ -ing then you ought to Ψ .” We will call this sentence “MEP”. On our analysis, this sentence is just straightforwardly a true conditional claim.⁷ It is true because the consequent ‘you ought to Ψ ’ just means that $\ulcorner \text{to } \Psi \urcorner$ follows from your evaluative attitudes. The antecedent guarantees that your set of evaluative attitudes includes at least that you intend to ϕ and that you believe that Ψ -ing is necessary for ϕ -ing. Thus $\ulcorner \text{to } \Psi \urcorner$ does in fact follow from these premises. Therefore MEP is true. But we deny that it is normative. If the objector insists that it is normative because it is the kind of thing you can violate, we think that by a “violation” the objector must have in mind a situation where you fail to take the means to your ends. This doesn’t make the means-ends principle false (it can’t be false) but rather, the norm that you violated is just the first-order normative claim that you ought to go north (or whatever the means happened to be in the particular case).

⁷ While it is standard to write the means-ends principle this way, notice that we now have a sentence with free variables in it, which is problematic. One natural thought is that it is really a universally quantified sentence quantifying over actions.

As a third try, the objector might think that there is some kind of meta means-ends principle that really says something like “you ought to obey the means-ends principle.” Here, the idea is that we have to actually accept and be guided by the means-ends principle. But we argue that an agent need not accept any such principle to be justified in their practical deliberation. This new claim is now of the form $\ulcorner S \text{ ought to } \phi \urcorner$ where ϕ is obeying the MEP. It is not entirely clear what it could mean to “obey” a conditional sentence such as MEP, but for now we will take for granted that this makes sense or that a suitable replacement could be spelled out. As a sentence of the form $\ulcorner S \text{ ought to } \phi \urcorner$, according to our thoroughgoing constructivism, this claim is indeed normative and will have the usual analysis: S ought to obey MEP just in case there are elements in S’s motivational set that entail MEP-ing. Certainly, ordinary cases of practical deliberation do not require that the agent assent to such a principle in any way.

The mistake can be clearly seen in the theoretical domain. If $P \vDash Q$, then all P worlds are Q worlds, and we say that in this case, provided we believe P, we ought to believe Q. It would be absurd to demand that in reasoning the agent must also accept another principle that tells them *that when $P \vDash Q$, and we believe P, we ought to believe Q*. This is the central lesson in Lewis Carroll’s “What the Tortoise said to Achilles” (Carroll 1895). We believe that the demand for a supporting practical normative principle is equally misguided. Our attitudes support other attitudes according to the model-theoretic entailment relations. The entailment relations themselves are neither practical claims—claims that could appear as premises of a practical argument—nor are they normative claims which are second-order statements that become true when the entailment relation holds. There are true sentences about such relations such as the sentence that says, “If P is true and P entails Q then Q is also true,” and there are also sentences about theoretical reasoning such as “If you believe P, and you also believe that P entails Q, then

you ought to believe Q.” Both of these sentences are true. Neither of these sentences is normative and you don’t have to accept either of them in order for them to be relevant to the standards of correctness for theoretical reasoning. Similarly, the facts about the logic of practical deliberation and sentences that describe such facts are true whether or not we assent to them.

While we think we have now argued for a thoroughgoing metanormative constructivism, there is a persistent suspicion that constructivism must ultimately boil down to a form of either expressivism or realism. Street herself invited this sort of suspicion when she concedes that the constructivist “agrees with the expressivist that normative terms are used to express a state of mind that is different from ordinary belief” (Street 2010: 376), and Korsgaard does not help when she says that “expressivism, I believe, is like realism also true after all, and also in a way that makes it boring” (Korsgaard, 2003: 122). Assertions like these have led to understandable confusion. Thus, Lenman (2012) presses Street to explain just what state of mind these normative judgments are, and he plausibly concludes that they must express a noncognitive valuing state in which case constructivists are a kind of expressivist after all (Lenman 2012: 218). Gibbard (1999) and Ridge (2012) have come to similar conclusions.

Earlier we distinguished between practical judgments which have the form $\ulcorner \text{to } \phi \urcorner$ and normative judgments which have the form $\ulcorner S \text{ ought to } \phi \urcorner$. We said the former are first-order commitments of an agent while the latter are second-order commitments about our first-order commitments. This allows us to decisively distinguish between expressivist and constructivist semantics. For the expressivist, normative judgments are a kind of practical judgment. When a speaker, S, says $\ulcorner A \text{ ought to } \phi \urcorner$ she takes a practical stand on A’s ϕ -ing. It is an expression of a fundamentally desiderative act of the will. As such, it is the sort of proposition that could be an element of practical reasoning. For the constructivist, when S says that $\ulcorner A \text{ ought to } \phi \urcorner$ she is not

taking a stand on ϕ at all. Instead, she makes a claim about what A's evaluative attitudes entail. Since normative judgments are distinguished from practical judgments because of their logical and semantic properties, constructivism is decidedly different from expressivism.

The case for distinguishing constructivism from realism is a little more complicated. After all, there are significant terminological debates about which views are appropriately called "realist." As Street (2010: 370) points out, one way of carving up the metaethical landscape is to claim that realism is just the view that normative claims can be true or false and that some are true. On this extremely weak understanding, any version of constructivism would count as a realist view. But on another natural division, stance-independence is a necessary feature of realism and so any constructivism, ours included, is not a realist view (see e.g. Shafer-Landau 2003: 15). We have already pointed out that on our view, normative judgments are not practical judgments. Combined with stance-dependence, this separates constructivism from both expressivism and realism.

Our view is not the only view that has both of these features. Railton's ideal observer theory (Railton 1986) and Copp's society-centered theory (Copp 2001) do as well. Copp (2013) treats these views as constructivist, since they can be characterized as involving a hypothetical construction procedure. We agree with Street (2010) that focusing on the practical point of view and practical deliberation yields a better characterization of constructivism, and on this understanding, these views would not be constructivist. Whether or not these views are properly thought of as constructivist, they are important to examine since they feature as prime examples in Copp's (2013) argument that there are no real important philosophical distinctions between constructivism and realism.

Copp argues that these views should count as versions of moral naturalism, which is a realist view. Instead of realism requiring stance-independence, Copp suggests a list of commitments that any realist must have and names someone who is so committed a “basic realist” (Copp 2013: 120). Whether these views should count as constructivism, realism, both, or neither, our view can be clearly distinguished from them by focusing on what makes normative claims true. Metanormative constructivism is not a version of moral naturalism because we hold that normative facts could not be the objects of empirical investigation. On our view, normative facts are just facts about what is practically entailed from a certain set, namely, the agent’s evaluative attitudes. While these entailment facts are not metaphysically mysterious, they are not themselves the objects of empirical inquiry. To illustrate the difference, Street (2010: 372-374) draws an analogy to baseball. She imagines an event, e.g. Derek Jeter slides into second base before the ball is caught, and she asks what would make it true that Jeter is safe. According to naturalistic stance-dependent views like Railton’s and Copp’s, what makes Jeter safe has to do with questions of what an ideal umpire would say where ideal could be spelled out in some naturalistic way. But, Street argues, according to constructivist views like our own, what makes Jeter safe is that his being safe follows from facts about the event combined with the rules of baseball. And questions of what follows from the rules of baseball are not themselves the object of empirical investigation.

We agree with Street and conclude that since our version of constructivism lacks the stance-independence of standard realist positions and the naturalism of Railton’s and Copp’s views, it occupies a distinct place in the metaethical landscape.

Nothing we have said here constitutes a defense of a metanormative constructivist position as opposed to traditional alternatives such as realism and expressivism. What we have

done, however, is to provide what we think is the best version of such a metanormative constructivist theory. According to this version of constructivism, a normative claim of the form $\ulcorner S \text{ ought to } \phi \urcorner$ is true if $\ulcorner \text{to } \phi \urcorner$ is entailed from S's set of evaluative attitudes. This understanding heeds the core intuition of the constructivist that normative claims are made true because of our practical commitments as agents, and it reflects the centrality of practical deliberation to normative truth. We conclude that Ridge's challenge to constructivists has been met.

References:

Behrends, J. and J. DiPaolo. (2011). "Finlay and Schroeder on promoting a desire." *Journal of Ethics and Social Philosophy* 6, 1–6.

Broome, J. (2015). *Rationality Through Reasoning*. Wiley-Blackwell.

Carroll, L. (1895). "What the Tortoise Said to Achilles." *Mind* 4(14) 278-280.

Copp, D. (2001). "Realist-Expressivism: A Neglected Option for Moral Realism." *Social Philosophy and Policy* 18 (2): 1-43.

Copp D. (2013). "Is Constructivism an Alternative to Moral Realism?" in *Constructivism in Ethics*, ed C. Bagnoli (Cambridge: Cambridge University Press): 108-132.

Darwall, S., Gibbard, A., and Railton, P. (1992). "Toward Fin de Siècle Ethics: Some Trends," *Philosophical Review* 101: 115–89. Reprinted in *Moral Discourse and Practice: Some Philosophical Approaches*. Eds. Stephen Darwall, Allan Gibbard and Peter Railton. Oxford: Oxford UP, 1997. 3–47.

Finlay, S. (2006). "The Reasons that Matter." *The Australasian Journal of Philosophy* 84(1): 1-20.

Finlay, S. (2014). *Confusion of Tongues: A Theory of Normativity*. Oxford: Oxford University Press.

Gibbard, A. (1999). "Morality as Consistency in Living." *Ethics* 110: 140–164.

- Gibbard, A. (2003). *Thinking How to Live*. Cambridge: Harvard University Press.
- Hussain, N. and N. Shah (2006). "Misunderstanding Metaethics: Korsgaard's Rejection of Realism." *Oxford Studies in Metaethics* 1: 265–294.
- James, A. (2012). "Constructing Protagorean Objectivity," in *Constructivism in Practical Philosophy*, ed. Lenman and Shemmer (Oxford: Oxford University Press, 2012): 60-80.
- Korsgaard, C. (2003). "Realism and Constructivism in Twentieth-Century Moral Philosophy." *The Journal of Philosophical Research*, APA Centennial Supplement, Philosophy in America at the End of the Century, 2003: 99–122. Reprinted in Korsgaard, C. *The Constitution of Agency: Essays on Practical Reason and Moral Psychology*. Oxford: Oxford UP, 2008, 302–326.
- Lenman, J. (2012). "Expressivism and Constructivism," in *Constructivism in Practical Philosophy*, ed. Lenman and Shemmer (Oxford: Oxford University Press, 2012): 213-225.
- Rawls, J. (1999). "Kantian Constructivism in Moral Theory," in *Collected Papers*. Ed. Samuel Freeman. Cambridge, MA: Harvard UP, 1999. 303–358.^[L]_[SEP]
- Railton, P. (1986). "Moral Realism." *Philosophical Review* 95(2) 163-207.
- Ridge, M. (2012). "Kantian Constructivism: Something Old, Something New," in *Constructivism in Practical Philosophy*, ed. Lenman and Shemmer (Oxford: Oxford University Press, 2012): 138-159.
- Scanlon, T. (1998). *What We Owe to Each Other*. Harvard University Press.
- Schroeder, M. (2007). *Slaves of the Passions*. Oxford. Oxford University Press.
- Shafer-Landau, R. (2003). *Moral Realism: A Defence*. New York: Clarendon Press.
- Sharadin, N. (2015). "Problems for Pure Probabilism about Promotion (and a Disjunctive Alternative)." *Philosophical Studies* 172(5): 1371-86.
- Street, S. (2008). "Constructivism about Reasons." *Oxford Studies in Metaethics* 3: 208–245.
- Street, S. (2010). "What is Constructivism in Ethics and Metaethics?" *Philosophy Compass* 5: 363-384.
- Street, S. (2012). "Coming to Terms with Contingency: Humean Constructivism about Practical Reason," in *Constructivism in Practical Philosophy*, ed. Lenman and Shemmer (Oxford: Oxford University Press, 2012): 40-59.
- Vranas, P. B. M. (2008). New Foundations for Imperative Logic I: Logical Connectives, Consistency, and Quantifiers. *Noûs*, 42, 529–572.

Vranas, P. B. M. (2010). In Defense of Imperative Inference. *Journal of Philosophical Logic*, 39, 59–71.

Vranas, P. B. M. (2011). New Foundations for Imperative Logic: Pure Imperative Inference. *Mind*, 120, 369–446.

Vranas, P. B. M. (2016). New Foundations for Imperative Logic III: A General Definition of Argument Validity. *Synthese*, 193 (6), 1703-1753.

Wallace, R. J. (2012). “Constructivism about Normativity: Some Pitfalls,” in *Constructivism in Practical Philosophy*, ed. Lenman and Shemmer (Oxford: Oxford University Press, 2012): 18--39.