

HUMEAN CONSTRUCTIVISM AND DELIBERATIVE COHERENCE

by

Danilo Doche Linhares

A Thesis Submitted in
Partial Fulfillment of the
Requirements of the Degree of

Master of Arts
in Philosophy

at

The University of Wisconsin-Milwaukee

May 2019

ABSTRACT

HUMEAN CONSTRUCTIVISM AND DELIBERATIVE COHERENCE

by

Danilo Doche Linhares

The University of Wisconsin-Milwaukee, 2019
Under the Supervision of Professor Nataliya Palatnik

According to Humean constructivism in metaethics, there is no incoherence in holding that different agents should act on aims that are not co-possible. I will show that this commitment undermines Humean constructivists' own treatment of normative judgments, where these judgments are meant to function both as prescriptions and assertions of fact. When ideally coherent Humeans engage others in conversation, their claims about others' reasons to act function as imperatives rather than as assertions; conversely, when Humean reasoners think of those claims while deliberating on their own, they carry no prescriptive weight at all. In light of these issues, I propose that coherence in normative judgment should take into account the joint realizability of agents' aims. To act on reasons involves acting on aims the agent thinks are worth pursuing. And actions whose aims are in conflict cannot be successfully performed together. I argue that where aims conflict, so do the prescriptions for acting on them.

TABLE OF CONTENTS

ABSTRACT.....	ii
ACKNOWLEDGMENTS.....	iv
1. INTRODUCTION.....	1
2. HUMEAN CONSTRUCTIVISM ON COHERENCE AND CORRECTNESS.....	3
3. TALKING AND THINKING LIKE A HUMEAN.....	6
3.1. TALKING LIKE A HUMEAN.....	9
3.2. THINKING LIKE A HUMEAN.....	12
4. ESTABLISHING THE MORAL COMMON GROUND.....	14
4.1. TELEOLOGICAL REASONS FOR ACTION.....	18
4.2. THE PUBLICITY OF VALUES.....	22
4.3. HOW INTENTIONAL ENTAILMENTS ENABLE NORMATIVE THOUGHT AND TALK.....	28
5. CONCLUSION.....	30
WORKS CITED.....	32

ACKNOWLEDGEMENTS

The completion of this thesis would not have been possible without the assistance of my advisor, Dr. Nataliya Palatnik. I am extremely grateful for her many important suggestions, and for the encouragement and support she provided me along the way. I would also like to thank Dr. Peter van Elswyk, whose teaching on philosophy of language led me to one of the central insights of this piece. He was also very generous with his time and provided timely feedback at crucial junctures of the writing process, which I am immensely appreciative of. I would also like to thank Dr. Stan Husi for his critical advice at the early stages in the development of this thesis, and for sitting in on my thesis committee. Lastly, I would like to thank John Muller and Dr. Carly Bertrand, as well as members of our Writing Workshop (Dr. Blain Neufeld, Dennis Moore, Jason Dewitt, Jean Pierre Cordero Rojas, Jordan Martin, Joshua Vonderhaar, Steven Canet, Sofia Paz, Owen Forbes, and Taojie Wang) for reading early drafts of my thesis and for helping me work through my prose. My writing has come a long way thanks to you – even if it still has a long way to go.

“There is nothing either good or bad, but thinking makes it so.” – William Shakespeare, *Hamlet*, act II, sc. 2, p. 11.

“Thinking is just talking to yourself. And since meanings are public, talking is just thinking in the company of others.” – Christine Korsgaard, *Self-Constitution*, p. 197

1. Introduction

Humean metaethical constructivists accept that if an agent can coherently think that he should torture others for the sake of his own sadistic pleasure, the agent’s coherent thinking ‘makes it so’¹. Opponents have argued that this claim about the sadists’ reasons is tantamount to approving of sadism². In response, Humean constructivists deny that making judgments about other agents’ reasons involves endorsing or rejecting them³. Humeans claim we can still have good normative reasons for condemning and stopping the sadist from pursuing his aims, even if we think that these aims are what he has the most normative reason to pursue. Importantly, this response to the sadist generalizes to all cases where agents have aims which other agents disapprove of. According to the view, there is no incoherence in holding that ‘A should φ for the sake of X’ and ‘B should χ for the sake of Y’, where X and Y are not co-possible. Only judgments about what a same agent should do can be incoherent with one another – and even then, only on instrumental and logical grounds.

Here I will show that by limiting the scope of practical rationality in this way, Humean constructivists undermine their own treatment of normative judgments. According to them, normative judgments function both as prescriptions and statements of fact. But when an ideally coherent Humean agent engages others in conversation, her claims about others’ reasons function only as imperatives, and not as assertions of normative propositions. For example, when she tells another agent that they should help her with something, conversational dynamics block her from stating that they have a reason to help

¹ Street (2009: 292ff).

² E.g. Blackburn (1984: 218ff), Gibbard (1999: 145ff), Scanlon (1998: 55ff; 2014: 107), Streumer (2017: 97).

³ Street (2009: 293-294).

her: she at best issues a prescription bereft of factual claims about their reasons. Conversely, when Humean reasoners deliberate on their own, claims about other agents' reasons carry no prescriptive weight at all to them. A Humean reasoner thinks of them purely as statements of fact.

In light of these issues, I propose that coherent normative judgments should take into account the joint realizability of agents' aims. I defend the idea that to act on reasons involves acting on aims the agent thinks are worth pursuing. I will draw on this idea to show a sense in which reasons for action are public. I argue that where aims of action cannot be jointly realized, the reasons for acting on those aims conflict⁴.

For the purposes of this paper, I will engage primarily with Sharon Street's 'formalist metaethical constructivism' (2008, 2009, 2010, 2016), as I take it to be the most clearly articulated representative of Humean constructivism in the literature today. My discussion will proceed as follows. First, I will briefly outline the Humean constructivist project in meta-ethics in relation to its cognitivist and constitutivist roots. In doing so, I will distinguish the Humean constructivist's account of normative judgments from that of traditional Humean ethics, and from Kantian metaethical constructivism. Second, I will show some of the limitations of the Humean constructivist picture of practical rationality. I will argue that the discourse of agents who reason according to its standards does not settle questions about what reasons agents have, as assertions would, but rather serves *only* to influence interlocutors' behavior. I will also show that because judgments about other agents' reasons entail no prescriptive contents for anyone other than the agents in question, Humean practical reasoning robs those judgments of their prescriptive force. This gives us reason for rejecting that pure means-ends coherence is sufficient for rational coherence in normative judgment. I

⁴ It is worth noting that my account of coherence in normative judgment will not directly give us standards of *correctness* in the way the Humean constructivists' claim to. I limit my positive proposal to the claim that, where aims conflict, the set of judgments that prescribe acting on one of those aims can only be correct to the extent that judgments that favor acting on conflicting aims are incorrect. Offering a fuller picture of correctness in normative judgment is a much larger project, one for which I lack the space and the resources to pursue here.

then argue that only if agents' aims are jointly realizable can we coherently judge that these agents have reason to act in pursuit of them.

2. Humean constructivism on coherence and correctness

Humean constructivism in metaethics owes at least as much to modern Kantian theories of value as it does to traditional Humean ethics. The view is Kantian in that it gives pride of place to standards of reasoning in their account of correct normative judgment⁵. Contra Hume, correct judgment is not a matter of having the proper moral affections, but of *judging* in accordance with constitutive standards of practical reasoning. Unlike desires and beliefs, the attitude of judging that something is a reason to act has both conative and cognitive elements for Humean constructivists. According to Street, purely conative attitudes like desires do not yield metanormative standards of correctness for other attitudes of the same kind⁶. In turn, purely cognitive, belief-like attitudes are not essentially motivating in the way 'taking something to be a reason to act' is. Humean constructivists differ from their Kantian counterparts primarily in that they disagree on what the constitutive standards of practical reasoning consist in. Kantians claim that these standards entail commitments to particular normative judgments: namely, always taking humanity to be a valuable end⁷. Humean metaethical constructivists share Hume's skepticism about vindicating particular judgments through standards of reason alone⁸. But unlike Hume, constructivists do not seek to justify judgments based on the character or relations between agents' sentiments *simpliciter*. Instead, they argue that facts about what an agent has reason to do are constructed out of agents' normative judgments about themselves, given the logical and instrumental relations between them⁹.

⁵ In fact, the case which inspired my sadist – and which Street discussed as an objection to *her* view (2009, 292ff) – was originally raised by Alan Gibbard (1999) as part of a critique of Korsgaard's Kantian constructivism.

⁶ Street (2008: 27, 42-43).

⁷ Korsgaard, forthcoming.

⁸ Street (2016: 297, 332).

⁹ Street (2010: 367).

To illustrate with an example: consider an investor deliberating about whether she should keep her investment portfolio when some of the companies in it engage in practices she disapproves of. According to Humean constructivists, in order to determine what she has most reason to do, we have to first look at how she weighs her own values in connection to one another¹⁰. The judgments that ‘she should not profit off of the unsafe working conditions of others’ and that ‘she should invest in whatever companies maximize her profits’ can come into conflict in light of the non-normative facts which are relevant to them (whether investing in a certain company maximizes her profits; whether the company’s employees work in unsafe conditions). But the investor may also have a network of other values that inform each of these two judgments: maximizing her profits can be instrumental to having a standard of living she thinks is valuable to have; not profiting off of inhumane treatment of workers might be instrumental to her aim of acting in ways that would make her partner proud – where her partner deplores these practices. Each of these relations form a segment in the ‘total web of the valuer’s normative judgments’¹¹. The correctness of any judgment about the agent’s reasons is a matter of how well that judgment sits within the agent’s own web; of how close to reflective equilibrium making that judgment would leave the agent¹².

By carrying out this scrutiny of her values, we arrive not only at the coherent judgements she can make about her reasons, but at the facts about those reasons. Scrutiny is not just our best bet at making correct judgments: its outcome constitutes their correctness. Humean constructivists like Street hold that “facts about reasons [are reducible to] facts about what we *judge* or *take* to be reasons”¹³; or, in more detail,

¹⁰ I am borrowing this example from Street (2008, fn. 38).

¹¹ Street (2008: 32).

¹² In Street’s words, the degree to which something is a reason for an agent “is a function of how strongly she holds the normative judgments in question and how close to the center of her total web of normative judgments they lie” (2008: 32).

¹³ Street (2008: 40).

reducible to “facts about what is entailed by [our normative judgments, given] the ‘rules of practical reason’ in combination with the non-normative facts”¹⁴.

For Street, these ‘rules of practical reason’ are given by the inferences we characteristically draw from practical reasons claims when interpreting them as such. To use one of her own examples, consider the following judgments:

- (1) “I have conclusive reason to get to Rome immediately.”
- (2) “Getting on a plane is the only way to do so.”
- (3) ‘I have conclusive reason to get on a plane’¹⁵.

If someone asserts (1) and (2) in full consciousness of what she is asserting, she ought to infer (3). Were she to infer otherwise (e.g. that “I have no reason to get on a plane”), we would not be able to make sense of her as making normative judgments at all by (1) or (3). Inferring (3) from (1) and (2) is required of her by the constitutive formal features of normative judgments. In that sense, judgment (3) is *constitutively entailed* by the conjunction of judgment (1) and the instrumentally relevant, non-normative fact expressed by (2).

Whether or not a normative judgment is correct is a function of whether or not they are constitutively entailed by other judgments. There is nevertheless a question of which other judgments to take into account. Meta-ethical constructivists do not assume any foundational values. Given a judgment about the reasons of a particular agent, the ‘other judgments’ are just whatever reasons the agent takes herself to have¹⁶. What is more, the correctness of all normative judgments gets established that way.

¹⁴ Street (2010: 373). The meta-ethical project of constructivism more generally is to offer a theory of moral facts that follows exclusively from considerations about what is entailed from within the practical point of view of agents, without taking for granted any substantive normative judgments. This is true of a number of constructivists across the board, including Kantians like Christine Korsgaard (1996, 2009) and Onora O’Neill (1989), neo-Aristotelians as LeBar (2008), among others, e.g. Velleman (2009).

¹⁵ Street (2008: 24).

¹⁶ To stress, for Humean constructivists, taking oneself to have reasons involves having a *cognitive attitude* towards a propositional content. This is an important point of contrast with how traditional Humean ethics thinks of moral judgment. Most prominent interpretations of Hume on the nature of moral judgment fall squarely in the

Grounding judgments like (1), which form part of the antecedent of some constitutive entailments, must themselves be submitted to the same sort of procedure. Facts about agents' reasons are constructed in cycle of cross-valuations where all of an agent's judgments are evaluated by the lights of what is instrumental to successfully acting on their other judgments. This makes the account avowedly circular: Humean constructivism offers a coherentist account of correctness in normative judgment. If a judgment would withstand scrutiny in reflective equilibrium, that means the judgment is correct¹⁷. (For the purposes of this paper, I am not interested in whether coherentism offers the right picture of correct practical reasoning. What I wish to dispute in the following sections is rather the picture of *coherent* practical reasoning offered by Humean constructivism.)

In summary, Humean constructivism tells us that if an agent judges that there is reason to act (or not to act), and this act or omission would hinder the attainment of some other aim the agent judges *she* should pursue, that counts against the judgment's being correct. Conversely, if the act or omission contributes towards the attainment of some aim the agent endorses *for herself*, that counts in favor of the judgment. Yet, facts about what other agents should do tell us nothing about what we ourselves should do, and vice-versa. This, I will show, creates some problems for the Humean pictures of normative discourse and deliberation.

3. Talking and thinking like a Humean

When we make judgments about what we have reason to do, we are not neutral about whether or not we should succeed in acting on those reasons. This foundational feature of practical reason explains why means-ends coherence figures among its constitutive standards. The judgments that 'I have reason

non-cognitivist camp: to the extent that propositional contents are involved at all, normative judgments *express* propositions about the affect-laden, evaluative attitudes of the spectator towards the object of evaluation (Mackie, 2003, Blackburn, 1993). For Hume, moral judgments are not *assertions* involving irreducibly normative predicates, such as the 'reasons' predicate, as Street would have it.

¹⁷ Street (2008: 36).

to get to Rome immediately' and 'getting on a plane is the only way to do so' *instrumentally* entail that 'I have reason to get on a plane' precisely because the agent's judgment commits her to procuring whatever means are necessary to realizing the aims which inform her reasons to get to Rome immediately. This relationship between aims and agents' reasons is understated in Street's original example. We can imagine that the agent does not take 'getting to Rome immediately' to be worth doing for its own sake, but that she values being in Rome so she can be with a loved one who suffered a serious accident. Being with her loved one as soon as possible serves as a common reason-giving purpose for both her getting on a plane *and* for her getting to Rome immediately: both acts are means (albeit at different levels of instrumentality) to a single aim.

The roles of agents' aims in normative judgment go beyond what the Humean lets on. I propose that an agent who holds the first two propositions in full consciousness also cannot say that "other agents have no reason to allow me to get on a plane". If she did, it would be difficult to see her as making normative judgments at all. Whatever she intends to achieve by getting to Rome immediately, she must represent it as a valuable end if she really takes it to give her reason to get there. From the point of view of that agent, other agents have at least *a* reason to cooperate with her getting on the plane: namely, those same considerations in favor of the aim she seeks to realize by getting to Rome. Much like with Humean means-ends coherence, I argue it is also constitutive of practical rationality that she cannot hold in full consciousness that "being with my injured loved one as soon as possible gives me conclusive reason to get to Rome" and that "other agents have conclusive reason to act in such a way that would prevent me from being with my loved one as soon as possible". The judgment that 'X gives A reason to φ ' *intentionally* entails that all agents have reason to be minimally cooperative with A's acting on the intentions informed by X¹⁸.

¹⁸ A couple of things to note about this discussion. First, I kept to Street's original formula of 'X gives A reason to φ ' (2008: 19). But the language of 'reasons' is inessential here. The same general formula could easily be extended to more idiomatic sentences containing expressions like 'must' (e.g. "I must get to Rome immediately to be with my loved one..."). Second, so long as we are *not* speaking of 'conclusive reasons' or other expressions for deontic necessity, the entailed normative judgments about other agents predicate merely *pro tanto* reasons, not conclusive ones. Our Rome traveler might, for instance, not have a ticket (or the money to buy one), giving

Humeans like Street will deny that intentional entailments set constitutive standards for practical reasoning in the ways that instrumental entailments do. I find them highly intuitive, but appealing to these intuitions alone is not enough¹⁹. Taking some inspiration from Korsgaard (2009), I will show that if we do not treat intentional entailments as constitutive of coherent practical reasoning, we end up with ideally coherent agents:

(3.1) who cannot settle any joint inquiries into normative matters, as they are largely incapable of adding normative propositions to the common ground of a conversation;

(3.2) whose own intrapersonal deliberation fails to assign prescriptive contents to judgments about other agents' reasons.

Consider someone whose car has broken down in a remote country road (call her Susan), and who needs to get help from another driver to repair it. Naturally, she takes this predicament to give her good reason to wave towards the first car passing by. She recognizes the first driver coming towards her: it is her selfish neighbor Joe. Being a good instrumental reasoner who is well acquainted with Joe's evaluative attitudes, she recognizes that it is instrumentally entailed from Joe's present normative judgments that he does not have the most reason to help her. Suppose that Susan reasons according to Humean standards of correctness. Given the information available to her, she will always make judgments about what agents have the most reason to do given what she thinks is instrumentally and logically entailed from within their practical point of view. I want to think about what Susan would be telling Joe if she told him that he should

airline employees a stronger reason for not letting him on the plane – reasons which that defeat whatever reasons the value of the traveler's 'being with his family' gives to airline employees. So, a penniless traveler who judges that agents involved in the airline business only have good reason to let her get on a plane should she pay for her ticket (which she cannot afford) could conclude that the right thing for her to do might not be getting on a plane after all.

¹⁹A Humean could argue that we are inclined to mistakenly extend our values to agents to whom they do not belong, as if normative qualities were *inherent* properties of the ends we intend to achieve and not qualities constructed out of particular agents' evaluative attitudes. Street argues to similar effect in (2008: 38).

stop and help her. I also want to think about what Susan would be telling herself by thinking that he should not.

3.1. Talking like a Humean

If Susan is thinking like a Humean, she thinks that she has good normative reason to ask Joe for help. Whether she thinks stopping to help is the best course of action for Joe or not, she may still think that the best course of action for her is to get him to help her. And doing that may involve telling him he has reason to help even if she thinks otherwise²⁰.

Regardless of whether Susan's claim on Joe reflects what she judges he should do, she is incapable of engaging him in joint normative inquiry. To elucidate this point, let me draw on Stalnaker's model of conversation and the pragmatics of assertion (2002). The model is not wholly uncontroversial. To avoid unnecessary disputes, I will assume only three basic tenets inspired by the view: (I) conversations are backgrounded by a set of presupposed propositions, called the common ground; (II) in the context of a conversation, propositions are part of its common ground if and only if all participants are disposed to behave as if they accepted them (that is their characteristic feature as conversational presuppositions)²¹; (III) the essential effect of an assertion is to update the common ground by introducing the asserted

²⁰ A central part of Street's response to ideally coherent immoralists is that morally upright Humeans have good reason to persuade others to do what is moral even if they judge that *should not* do so by their own lights. (See Street, 2009: 293-294.) Street's response is in similar spirit to the consequentialist pragmatics Korsgaard argues against as she sets up her own view of the aims of interaction: "For Sidgwick, talking to someone is always a matter of trying to do things to him, (...) to cause him to have certain ideas. So Sidgwick can't see why we should count the immediate effect of using a word as telling the truth, while the more distant effects of using it - the inferences drawn and so forth - count as further consequences. For Sidgwick, it's *all* just so much spin." (2009: 197). My objections against the way Humean constructivists handle disagreement among coherent agents parallels Korsgaard's treatment of Sidgwick: a lack of publicity renders interlocutors incapable of deliberating together in conversation.

²¹ Following Stalnaker, I treat 'acceptance' as the broad category of cognitive attitudes involved in treating a proposition as correct or true (2002: 716).

proposition to it: an assertion is successful if participants become disposed to behave as if they accepted its contents²².

When speaking of dispositions, I do not mean to suggest that accepting a proposition is always *directly* motivating. Rather, when a proposition is added to the common ground, conversation participants will be disposed to take its correctness for granted should the occasion arise. That may have repercussions for participants' speech-acts as well as for their non-verbal behavior. Consider updating the common ground with the proposition 'there are children sleeping upstairs'. On the verbal end, once the proposition is added to the common ground, participants may use definite noun-phrases (e.g. 'the sleeping children upstairs') to refer to the objects of quantification in the original proposition. On the non-verbal end, participants may refrain from making loud noises while in the area. The emphasis on dispositions tells us, for one, that participants need not *engage* in any particular behaviors in order to presuppose the proposition²³. Second – and most importantly for our purposes – they also need not *actually* accept it. A proposition can be part of the conversational common ground even if one or more participants do not take it to be true, so long as those participants are still disposed to behave as if they did in conversation.

It is part of the Humean picture that agents' normative judgments are directly motivating, at least when those judgments are about the agents themselves²⁴. With this in mind, let us return to Susan and Joe. Imagine she sees him coming in the distance and calls his cellphone. She then asserts: "you should stop and help me". Susan asserts this in an attempt to make this proposition part of the common ground of the

²² Street conceives of normative judgments as having both cognitive and conative components: making one *includes* having an attitude of acceptance towards a normative proposition, in addition to being at least somewhat motivated to act (2008: fn. 37).

²³ E.g. they need not use the definite noun-phrases at all (they can refrain from talking about the sleeping children upstairs and still presuppose that the proposition is true).

²⁴ This motivation may not be "strong enough to result in action in every case". But motivational obstacles notwithstanding, making a normative judgment disposes one to act as prescribed (Street 2008: fn. 37).

conversation between her and Joe, making Joe disposed to behave as if he accepted it. For Joe, being so disposed involves being disposed to stop for her.

Susan does not speak as she does idly. When she asserts that Joe should help her, she is acting from her own reasons. And the normative judgment that motivates her to make that assertion is her judgment that ‘she should enlist Joe’s help in fixing her car’. This motivation is independent of whether or not she judges that ‘Joe should stop and help her’. The fact that she actually judges that Joe should *not* stop does not dispose her to behave in any particular way. If she behaves as an ideally coherent Humean agent would, nothing but her own values should impact her motives²⁵.

The judgments of an ideally coherent Susan about what Joe should do don’t even dispose her to assert them in conversation with him. When Susan asserts that “you (Joe) should stop and help me”, she is *not* acting on a disposition she would have if she accepted it. Rather, she is acting only as if she accepted that ‘she should enlist Joe’s help in fixing her car’. Given what has been said about conversational presuppositions in (I) and (II), Susan’s assertion does not update the conversation’s common ground with the proposition she asserted, regardless of what Joe is disposed to do in response. All participants have to be disposed to behave as if they accepted the proposition. For Susan’s judgment about Joe’s reasons, there are no dispositions that fit that description: since Susan is disposed to act as correct Humean practical reason prescribes, only her normative judgments about herself dispose her to act one way or another.

²⁵ We can of course imagine that Susan values honesty more highly than enlisting the aid of the first driver (Joe) to come by. But if that were the case, this would put her judgments that ‘she should enlist Joe’s help’ in *instrumental* tension with the conjunction of her judgments that ‘Joe should not help her’ and that ‘she should get others to accept only propositions she herself takes to be true’: the only way for her to get Joe to help her then would be by deceiving him with respect to the first conjunct, thus acting against what the second conjunct tells her she should do. To make what she is disposed to assert dependent on what she judges Joe should do this way only reinforces the tension between judgments with conflicting aims (which is a tension Humeans have to deny). The same sort of reasoning applies to the possibility that Susan is a highly sympathetic agent and values Joe’s ability to do what she judges he has most reasons to do (by Humeans’ instrumentalist standards).

It becomes exceedingly difficult for Susan to introduce normative propositions about her interlocutors to the conversational common ground. Even if she and Joe were to hold the same judgments about Joe's reasons (e.g. 'Susan should give Joe some money in order to recruit his help' and 'Joe should help Susan in order to get a payment from her'), the latter does not inform Susan's conversational dispositions – just like the former would not inform Joe's if he too were an ideal Humean reasoner. No matter what Susan says, she is always first and foremost talking about herself. (We can imagine that Joe knows about Susan's conversational dispositions²⁶. If he does, it is hard to see how she could be communicating anything about *his* reasons when she makes normative claims about him – even if we extended the scope of our pragmatics beyond that of assertions.) If we see all joint inquiry as a matter of introducing common presuppositions to the common ground, it follows that there can be no inquiry into normative matters among ideal Humean reasoners.

3.2. Thinking like a Humean

Susan's judgments about other agents' reasons are dispositionally idle in conversation. This is relevant for what we can say about her ability to make normative judgments about other agents' reasons when thinking to herself. It is hard to see what is normative about them when they have no uniquely prescriptive role in Susan's individual inquiry about others' reasons to act. Like with talk, Susan can only think in prescriptive terms when she is thinking about her own reasons to act.

In the non-reductionist semantics we are working with, normative judgments cannot be reduced to judgments *about* what is entailed by agents' other normative judgments²⁷. This means that judgments about what is instrumentally and logically entailed from within one's practical point of view are not themselves normative judgments. An agent can still conceive of a normative judgment that is incorrect by

²⁶ The ability to tell how interlocutors' verbal behavior relates to their presuppositions is key to avoid conversation failures, where participants are mistaken about what their interlocutors presuppose.

²⁷ See Dorsey (2012).

Humean standards as giving her prescriptions about what to do (prescriptions she, on reflection, may think she should not follow)²⁸. Conversely, simply believing that certain normative judgments are instrumentally entailed by her other first-personal normative judgments does not yield any prescriptions unless she reasons according to the standards of coherence which Humean constructivists prescribe. That is: the claim that ‘X gives her a reason to ϕ ’, and the claim that ‘X gives her a reason to ϕ ’ is instrumentally entailed from within her practical point of view’ do not mean the same thing to her²⁹. The former only follows from the latter if she accepts that she *should* act on normative judgments instrumentally entailed by the set of her other judgments. Judgments about what is entailed from her practical point of view have to be mediated by a metanormative creed in the Humean construction procedure in order to have prescriptive force for her. Accepting that creed, according to the Humean, is the mark of coherent practical reasoning.

If Susan reasons like a Humean, making judgments about other agents’ reasons would only tell her something prescriptive in light of some independently held judgment about what her own reasons are. Because Susan’s judgments about what others should do have to be mediated by judgments about herself, the way they bear on her dispositions is functionally identical to the way non-normative judgments do – such as judgments merely about ‘what is entailed’ from hers or others’ practical point of view.

Let’s say Susan judges that ‘Joe should always wear Oxford shoes to work in order to impress his boss’, where this is a claim about Joe's normative reasons to wear Oxford shoes to work. Suppose also that she judges that ‘Joe would always wear Oxford shoes to work in order to impress his boss’, where this is a non-normative modal claim about Joe's dispositions involving wearing Oxford shoes to work. Her judgment about Joe’s reasons do not figure in her deliberations any differently than her non-normative one: she could judge that she should prevent him from acting as she thinks he would, in just the same way

²⁸ E.g. she understands “I should not ask Joe for his help” as having both prescriptive content (‘refrain from asking Joe for help’) and descriptive content (‘I have most normative reason not to ask Joe for help’), even if she thinks the judgment is incorrect.

²⁹ Treating these two claims as synonymous leads to a grounding regress, as shown by Dale Dorsey (2012).

she can judge that she should prevent him from acting as she thinks he *should*. One important difference between the two judgments is that the former is grounded by facts about what is entailed from within his practical standpoint. But that alone adds nothing of prescriptive value to her judgment. After all, even when she is thinking about what her own reasons are, judgments about what is entailed have no immediate normative force. Since they get no prescriptive force by way of how they are grounded, Susan's judgments about Joe's reasons are limited to playing the same indirect roles in her practical reasoning as any non-normative judgment of hers. Since Humean practical reasoning does not posit any formal connection between her judgments about what Joe should do and her judgments about herself, the former give her no *prima facie* reason to care about whether he acts as she thinks he should, any more than she has reason to care about whether he acts as she thinks he would. Thinking like a Humean robs Susan's judgments about other agents' reasons of their prescriptive character.

4. Establishing the moral common ground

Under the Humean picture of practical reasoning, an ideal agent has a lot of trouble talking to other agents about what they should do: she cannot introduce propositions about how they should act into the conversational common ground. Relatedly, ideal Humean agents do not represent other agents' reasons as having anything of inherent value to them, any more than they might represent any non-normative qualities as valuable. Her evaluative attitude towards other agents' reasons is *prima facie* value-neutral: they do not tell her anything normative unless she thinks of them under the lights of her own judgments about what *she* should do. It is a major strike against the view that its ideal agents struggle with moral discourse and thought in this way.

Humean constructivists claim that preserving instrumental entailments is both necessary and sufficient for coherent practical reasoning. As metaethical coherentists, Humean constructivists also believe that coherence is necessary and sufficient for *correct* practical reasoning. I do not intend to take a

stance on coherentism in this paper. I will discuss standards of deliberative coherence as a self-standing issue, apart from the nature of the relationship between coherence and correctness. What I hope to have shown above is that instrumental entailments are not sufficient for coherence (though they are still necessary).

I will now argue that including a new kind of constitutive entailment as a requirement for coherence will yield a better picture of the pragmatics of normative assertions, and of the prescriptive force of our judgments about the reasons of others. I have already gestured at what those entailments might look like: they connect judgments about agents' reasons with those about the reasons of others. I called them intentional entailments, and suggested that they have to do with the joint realizability of agents' valued aims. I can now flesh out this proposal a bit more, showing how it addresses some of the limitations of Humean standards.

Street treats 'having reasons' as a three-place relational property between a reason (*qua* unary predicate), an act, and an agent, following the general formula of "X is a reason to φ for agent A"³⁰. In order for something to enter in this relation as a reason, it has to be a consideration that can supply an agent with an aim for their action. Taken as reasons, the likes of 'being with one's injured loved one' or "the fact that a policy would benefit the wealthy at the expense of the poor"³¹ inform the agent of which ends are worth promoting (or preventing) through action: e.g. 'that one is with their loved one soon', 'that the policy is not implemented'. To say that something provides someone with reasons to act involves³² saying that some aim is valuable for her.

³⁰ See Street (2008: 19). For a more detailed treatment of how reasons might function as a three-place predicate that accords to this Scanlonian formula, see also Bedke (2010).

³¹ Street (2008: fn. 7), after Rawls.

³² My proposal does not require taking a stance on whether saying that something is valuable can be reduced to claims about reasons, or vice-versa. Here and in the following sections, this is what I intend to signal when I use the word 'involve', rather than 'entail' or 'is entailed by'. Judgments about what the normative reasons are (*qua* unary predicates) may be *grounded by* the correctness of judgments about which aims are most valuable, or it may

The case for my positive proposal relies on an Aristotelian notion of intentional action which I won't have the space to defend in this piece. It is the idea that all intentional action aims at promoting some ends³³. This is true regardless of the nature of the agent's reasons to act. Even if we concede that not all considerations in favor of doing something are claims about what ends are worth promoting³⁴, any agent who acts on reasons is still aiming at the realization of some end insofar as we can say that she is acting intentionally. When appropriate, we might broaden this idea about action to encompass intentional omissions as well. It seems plausible to think that intentional omissions also aim at something even though they are not actions. These aims are what distinguish omission from inaction through neglect, in much the same ways they distinguish actions from other non-intentional acts.

Practical aims, as ends of action, are also conceived along Aristotelian lines. An action or omission³⁵ might have as its aim something that is more like a static product: a feeling (e.g. pleasure), a fabricated object (e.g. a beautiful work of art), an event or the avoidance of one (e.g. saving a child from drowning), etc. I will refer to these ends as 'products', and the aims that target them as 'poietic aims'. I also want to say that action can aim at realizing a process, where the activity of the agent herself is the end the agent is trying to promote: these are activities like being a good friend, being courageous, competing or dancing for its own sake³⁶, etc. I call these activities 'praxes', and the aims that target them 'practical aims'. It is

be the other way around. Whether we prefer to characterize it one way or another will depend on which side we take on the debate between reasons- and values-fundamentalism. Street follows Scanlon in taking reasons as the fundamental normative notion (2008: 37), but this is only important to her theory insofar as she offers an account of *correctness* in judgment. For her, correct reasoning according to a standard of coherence is what grounds judgments about value.

I take no stance on such issues. All that is required for my purposes – and which I will discuss in the remainder of this piece – is that teleological considerations always play a role in determining our reasons for action.

³³ I am borrowing this analysis from Korsgaard (see 2009: 9-11; 101ff).

³⁴ Scanlon for instance argues that some reasons for action are not teleological in this way (1998: 89ff).

³⁵ Henceforth, when speaking of 'action' I won't distinguish between features of action and omission unless otherwise noted.

³⁶ 'For the sheer joy of it' as opposed to doing it for the sake of some product (like pleasure), as Korsgaard puts it (2009: 12, 121).

important to take note of these second kinds of ends and aims, as Scanlonian challenges to the teleological nature of reasons are founded on a neglect for their existence.

Like Street, I will argue for my proposal based on claims about what is constitutively involved in judging that someone has reasons to act. I will argue that all reasons for action either consist in, or entail some reasons to act which are both *teleological* and *public*. These two features of reasons are intimately connected. Ends of action set public conditions for their success – where success in action is something we are committed to in practice whenever we act intentionally. Insofar as we recognize that this practical commitment is shared by all agents, we are rationally required to understand ourselves as making claims on all agents whenever we talk about any one agent's reasons for action.

I will discuss each one of these two features of intentional entailments separately.

(4.1) Intentional entailments invoke a teleological conception of reasons.

In particular, I wish to maintain that all existential reasons claims of the form 'A has some reason to φ ' entail an evaluative claim about the aim of A's φ -ing (some end E). This is the claim that, from A's practical point of view, E is worth promoting. This position should be uncontroversial for a Humean, but as we will see, there are some strong arguments against it that need to be addressed.

(4.2) Intentional entailments also require that the value ascribed to the ends of intentional action are public, at least in a weak sense.

What I propose is that recognizing that some end E is worth promoting from an agent's point of view involves recognizing a *pro tanto* reason not to act in ways that prevent E from being realized. The converse is true for when agents' aims are negative poietic aims: if an agent's choice aims at avoiding some state of affairs, avoidance of that state of affairs is her end F. The reasons for that choice therefore entail that all agents have *pro tanto* reasons

not to act in ways that promote the state of affairs. We can say in those cases that E and F have the properties of being (publicly) valuable.

4.1. Teleological reasons for action

To restate a point from section 3: the central intuition my proposal shares with Humean constructivism is that agents are never neutral with respect to whether or not they succeed in acting on their own reasons. Success, in turn, is a measure of how well agents realize the aims of their actions – whether these aims are for practical or poetic ends. Thus, to say of some agent that she has reasons to act is to say that she has reason to act in ways that realize the ends at which her actions aim. What I propose in this section is that, by way of aiming at realizing those ends, the agent takes those aims to be worth promoting *for her*. In the next section, I will discuss the idea that she also takes them to be valuable *for everyone*. I will also turn to the relationship between these two properties of ends – ‘being worth promoting’ and ‘being valuable’. I do not wish to treat them as equivalent.

It may seem obvious that reasons for action have a teleological structure. But there is a strong case to be made against the idea that the existence of reasons for action entail that the action’s ends are to be promoted. One prominent example is offered by Tim Scanlon in *What we owe to each other* (1998). There, he argues against what he calls a ‘teleological conception of reasons’³⁷ which has much broader scope than the one I intend to defend here. In this section, I want to briefly discuss the part of Scanlon’s arguments that touch directly on what my intentional entailments require of the notion of reasons. I hope to do this without needing to settle other debates Scanlon is engaged with in those passages. Particularly, I do not wish to take a stance on whether reasons are fundamental over values³⁸, nor do I wish to take a stance on

³⁷ 1998: 80ff.

³⁸ My proposal is compatible with the claim that all reasons are grounded by a contractualist procedure involving their reasonable rejectability, as Scanlon proposes (1998: 153ff). Under this grounding framework, it would still be only by virtue of the fact that there are reasons to act on certain aims that we can say that those aims are non-instrumentally valuable. According to my proposal, it would just be the case that all of those reasons to act either

whether or not there are any reasons for action which are not teleological reasons. What I wish to maintain, contra Scanlon, is that whenever agents have reasons to act, at least some of those reasons are teleological reasons (as defined in item 4.1 above).

Scanlon argues that agents can have reasons to act, where none of these reasons will involve a judgment that the ends of her action are worth promoting, even if only for her. To show this, he presents action from deontological principles as a kind of counterexample to the idea that reasons for action must involve evaluative judgments about ends. Scanlon asks us to consider a principle which states that ‘saving many lives does not justify intentionally taking the one’. If Scanlon’s analysis is right, an agent who accepts this principle and takes it to be action-guiding will treat it as a reason in her judgment that ‘she ought not to kill one person in order to save many’³⁹.

Scanlon is right to note that, given a reason like this, what the agent judges she should do is not a matter of comparing the different intrinsic (dis)values of murder and human lives – even if we conceive of these values in agent-relative terms. An agent who thinks that a deontological prohibition against killing gives her reasons to act (or to omit herself) does not think that ‘murdering’ is of greater disvalue than the *poietic* state of affairs where a greater number of people are killed, such that the agent has an obligation to minimize the former. If one’s principle is that ‘saving many lives does not justify intentionally taking the one’, what this principle gives the agent reason to do is to ‘not kill for the sake of saving a greater number of lives’. An agent who considers this principle to be valid as a reason would not take herself to have reason to prevent some utilitarian from pulling the proverbial lever for the trolley tracks, for instance. It does not give the agent reason to minimize ‘killing for the sake of saving a greater number of lives’.

are or entail a teleological reason: reasons that predicate of an aim that it is to be promoted by the agent who has them.

³⁹ 1998: 84.

In fact, the very idea of comparing the values and disvalues of killing *vis.* saving lives goes against the spirit of the agent's reason *qua* principle: “[when a deontological principle is correct], one does not need to balance the value of abiding by it [e.g. saving lives] against the good to be achieved through its violation [e.g. not killing]”⁴⁰. According to Scanlon, what the principle states is precisely that considerations like these – that killing is more or less disvaluable for me than the saving of lives – are irrelevant to the question of whether or not an action (e.g. killing) is justified. That, Scanlon believes, shows that the principle is reason-giving independently of representing any ends as to-be-promoted. It shows that in some cases, such as when a valid deontological principle can be invoked, weighing the values of different ends of action can be treated as irrelevant in normative inquiry.

But Scanlon is wrong to reach this last conclusion. Even if we grant that agents who act from deontological reasons do not have reasons to maximize any particular state of affairs, that doesn't show that the agent does not attach a positive value to the goals she seeks to accomplish by her chosen acts or omissions. The agent's lack of commitment towards minimizing instances of killing is only establishes this point if we conceive of agents' ends exclusively as products. In the case of action for deontological reasons, it is clear that the criteria for success in action do not have to do with how many instances of a kind of state of affairs or property the agent is able to bring about. But that does not mean that there isn't a practical end involved in her reasons for acting.

When the agent chooses not to pull the proverbial lever of the trolley tracks, her omission is intentional. There is still something that the agent is trying to achieve by it, and that end is the measure of her success in action. That end is the form of her own activity: she aims at not killing (at least not from utilitarian motives), where *her* successful performance of the activity (or omission) is itself what she is trying to achieve by it. She could still fail if, due to weakness of will, she goes ahead with the killing in spite of

⁴⁰ Ibid.

intending the contrary. The badness of that failure in her eyes is on some level⁴¹ explained by the worth she attaches to her practical end. Not being a killer is something that she cares about, and that give her reasons to act in ways that promote that end, as a form of her own activity⁴².

What I have said so far still leaves open the question of whether or not she must compare the worth she attributes to her practical end ('not being a killer') to the unworthiness of the state of affairs where people are killed as a consequence of her omission. I suggested earlier that I think this is certainly not the case if we conceived of her end as a poetic one. But whether or not this also applies to the actual practical ends recommended by her deontological principle will depend on whether we want to say that bringing about that state of affairs is something she does intentionally. One possibility is that, if in the circumstances in which she chooses to act this state of affairs is necessary for her to succeed, and she knows this, then she brings it about intentionally. In that sense, bringing it about is one of her ends: she aims at it by virtue of choosing her omissive conduct, as a condition for her success in action. If intentionality and aims map onto each other as this sketch suggests, then she must be making a compared judgment about the value the aforementioned state of affairs as one of her poetic aims. On the other hand, if it is not appropriate to speak of this state of affairs as an aim of hers even in circumstances in which she knows bringing them about is necessary for her to act successfully, then the compared evaluative judgment is not rationally required.

⁴¹ That level may not be the most fundamental level of explanation. Again, I do not take a stance on whether 'the worth she attaches to her practical end' is itself explained by her reasons to choose the action which aims at those ends.

⁴² It is important to distinguish this claim from the idea that the agent's reasons not to kill in this case are agent-relative. We do not need to say, as Nagel (1989) does, that 'not killing' is worth having as an end only for her. Given what is required by my proposed intentional entailments, she could well judge that 'not killing' is a valid practical end for everyone, and still not take herself to have reasons to prevent anybody else from failing to act as they should as much as she can. What explains the departure from maximizing rationality is instead a change in the nature of the aims of action recommended by these principles. The ends which these principles cast in a positive light are usually practical rather than poetic. Whether her acts are such as to bring about that other people succeed in acting on their practical reasons is immaterial to her own success in acting on hers.

In summary, I hold that if an agent does something intentionally, she can't exclude from consideration the worth that she attributes to what her action is aiming at. Her action depends on the realization of its ends to be successful, and she is trying to succeed whenever she acts intentionally. This is one way to elaborate the Aristotelian conception of intentional action outlined above: that all intentional action aims at some end. Insofar as the agent treats her action as worth choosing by performing it intentionally, she is practically committed to the worth of its end.

It is worth noting that I am separating talk of 'value' from that of 'worth' here – where values are public but worth (or 'to-be-doneness') is not. This choice of terminology is somewhat arbitrary, but the distinction is not. I intend to use the notion of worth to capture the agent-relative component of valuations of ends, and to distinguish it from the kinds of public valuations of ends which I discuss under the label of 'value', which I briefly spell out in the next section.

4.2. The publicity of values

Since I have been discussing Scanlon, let me start by reiterating that the considerations about 'worth' in the previous section, and the ones about value to follow, are compatible with his buck-passing account. By way of saying that, I should be able to clarify the proposed distinction between value and to-be-doneness a little bit further. For our purposes, it is sufficient that for any end E, its being valuable entails that some agent has reasons to act in ways that promote E. This is something Scanlon explicitly accepts⁴³. It does not matter whether the entailment holds because E's value itself gives the agent reason to promote it (in which case it might give all agents equal reason to promote E), or if the normative buck is passed to other reason-giving properties of E. And that does not matter precisely because what is valuable need not overlap with what is worth pursuing (even though it sometimes might). I defend that if something is

⁴³ 1998: 95-98.

valuable, then all agents have reason not to act against. This is compatible with, but distinct from the Moorean thesis that it *also* follows from something's being valuable that agents have reason to promote it.

As shown in the previous section, to think that someone has reasons for action involves thinking that the aims of that action are worth pursuing for that agent. Here, I want to show that there is an additional necessary condition for thinking coherently about agent's reasons. In order to think that agents have reasons for action, we must think that the ends of those actions are valuable, in the restricted sense I proposed above. Being valuable involves only being such that *some* agent has reason to promote it (it is 'to-be-done' for someone), and that *all* agents have reason not to thwart efforts to promote it⁴⁴. The second conjunct on the claim defines the sense in which values are public. It gives us the ability to make universal claims on one another.

To lend plausibility to the claims above, this section will lay out a response to two related worries. First, why is it that agent-relative reasons for action (i.e. reasons involving agent-relative to-be-doneness) entail anything at all about other agents' reasons? Ultimately, my main point of contention with the Humean constructivist is that some entailments like this are part of our ideal of deliberative coherence. But to see why coherence in practical reasoning calls for thinking about the reasons of others in relation to ours requires some further elucidation of what has been discussed in section 3. Second, why is it that those reasons entail claims about what other agents ought *not* to do specifically? What is stopping me from claiming the 'value' of some ends gives everyone reasons to pursue it; or, as I have been using the terms here, that the choice-worthiness of certain ends holds for all agents, and not just some? As I have said in the beginning of this section, my proposal is neutral with respect to this second question. But my reasons for stopping short of this stronger claim also call for further discussion.

⁴⁴ I don't intend to suggest by this anything metaphysical about the existence of valuable ends. It is second nature to this project how or if these values are grounded.

Let us start from the first question. Believers in private reasons for action don't all believe that they are private to the same degree. Humean constructivists lie on an extreme of that spectrum. For them, not only are all reasons *prima facie* private, but fully so. Unless I make an evaluative judgment about my own reasons which expresses sympathies for others' success in doing what they should, their reasons for action tell me nothing about what I should or shouldn't do. For someone who is not a realist about reasons, and so does not believe that there are reasons out there in the world waiting to be discovered, it is fair to ask: why should somebody else's reasons tell me anything about what I should do? They are, after all, somebody else's reasons. There needs to be something about me as a rational agent – something about either the structure or the contents of my deliberative faculties – to vindicate this connection between the reasons of others and mine. The Humean mistakenly believes that only facts about the *contents* of one's evaluative judgments can bridge this gap between the first- and third-personal perspectives on reasons. There is something about the structure of coherent deliberation which the Humean is missing.

Think back to Susan, but from a first-personal point of view. That is: think of her as an agent like you, equally capable of acting for reasons and equally committed to successfully doing what you intend to do in practice. It is impossible for her to think that she should get help from somebody without thinking that this somebody should help her. Denying the latter, whether in thought or talk, would undermine her judgment that she should get help insofar as it is impossible for her to succeed at getting help if nobody helps her. Something similar would hold for our agent trying to get on a flight to Rome, in relation to another agent who is trying to stop her: she might concede that this other agent should in fact stop her, but she had better have good reason for it⁴⁵! So far, so obvious.

Now it shouldn't be too hard to switch to the third-personal point of view. Picture a coherent agent making judgments about her reasons, with no consideration about whether you yourself believe

⁴⁵ For more on this, see the discussion in fn. 18 at the beginning of section 3.

these judgments are correct. Put in general terms: if from her point of view an agent judges she has reason to φ , and we think of her as a coherent thinker, we will attribute to her the judgment that *nobody* ought to act in ways that prevent her from realizing the aims of her φ -ing – at least not unless they have a good, defeating reason for it. Coherence in practical reasoning, for the Humean as well as for me, trades on agents' ability to judge what best allows them to succeed in acting for reasons. We therefore attribute that judgment to our agent because successfully doing what she intends to do depends on her not being stopped. (It depends on other things as well, some of which may involve other persons. We will turn to these shortly, when we discuss the second question.)

Only one more imaginative move is required to conclude our elucidation on the first question. Imagine now *not* that you have a belief about what another agent judges her reasons to be, but that you yourself are making a judgment about what she has reason to do. If you make this judgment in the same way the agents in the previous steps have theirs, it won't necessarily have much to do with what the other agent cares about. The agents in the previous steps did not need to characterize those other agents at all in order to make claims on them. What is more: if this judgment is practical to you, and not just a judgment about some psychological features of the other agent⁴⁶, it ought to hold sway from *your* practical point of view as well. If these judgments about others are practical to you, they must involve some commitment to agents' success in acting for the reasons you believe they have. If your judgment has no connection to success in action, it is difficult to find the sense in which it is practical. This is the same line of thinking that I developed earlier, in section 3.2 of this essay: commitment to other agents' success in action is necessary to avoid the psychologism of the Humean account of normative judgment. We can now address the

⁴⁶ If some version of Cornell realism is right, it may still *follow* from such psychological facts. My proposed standard of deliberative coherence is neutral with respect to that.

second question, and spell out what it means, from the first-personal point of view, to judge that some other agent has reason to act.

I propose that we can generalize the structure of the practical commitment outlined above⁴⁷, such that making judgments about others involves attributing to oneself (and one's fellow agents) the judgment that one ought not to act in ways that prevent the other agent successfully doing what you believe she should. This negative requirement, rather than a demand for positively contributing towards other agents' success, is the minimum I believe we need in order to be able to make sense of ourselves as practically committed to the success of others in doing what they should. What I need to show is why coherence under intentional entailments does not require that reasons be public in the stronger sense. I do wish to argue that they are or are not required according to some theory of correctness in normative judgment.

Let us again look at things from the first-personal point of view. If I need others for doing something I should do (e.g. fixing my car – think of the Susan example again), then my reasons for action do allow me to make positive claims on others. The distinction between agent-relative to-be-doneness and public values is not one between considerations in favor of action versus considerations in favor of certain kinds of omissions. It is the distinction between actively working towards a given end versus only doing that which is necessary *not* to prevent that end from being realized. The latter may include both refraining from certain actions they would otherwise choose as well as engaging in activities which they ought to perform in light of their special importance for the realization of some end in the circumstances. If I am Susan, what kind of conduct is required for other agents to count as being committed to my success in getting my car fixed will depend on the extent to which their conduct is necessary for my success. If I can

⁴⁷ To say that we generalize the structure above to all agents is to say that this reasoning is necessarily *valid* for all agents, but not necessarily *sound* for all of them. Since I am not taking a stance on coherentism, for my purposes agents may be perfectly coherent and still wrong about practical reasons. I might thus disagree with an agent who thinks that they have a reason to φ . But if I judge that they do have reason to φ , then I must also think that I have a reason not to act in ways that thwart the agent's success in φ -ing. The correctness of anybody's thinking is immaterial to the question of coherence here.

get the car fixed as intended without anybody's assistance, then it is not necessary for my success that anybody help me. Since fixing the car is worth choosing for me, other agents who endorse that judgment can expect me to succeed so long as I am able to attempt the act. They can therefore be invested in my success so long as their intentional actions are consistent with what they believe I need in order to succeed. To put the point negatively: it only becomes difficult to see them as caring about whether or not I succeed to the extent that their actions work against my success in their eyes.

To summarize: in order to avoid the issues in joint and individual deliberation I presented in section 3, I propose that ideal reasoners should think of these valuable ends as public in at least a narrow sense. Deliberative rationality involves representing the set of all agents' valuable ends as a single, *consistent* set of ways in which the world should be. These aims may be valuable for particular agents in different ways. E.g. if "my summiting the Kilimanjaro" gives *me* a *pro tanto* reason to do what it takes to climb successfully, that valuable end need not give everyone a reason to do what it takes to promote my climbing. What the demand for a consistent superset of all agents' valuable ends does yield is that all agents have a *pro tanto* reason not to act on aims that conflict with my climbing. Judgments of the form "realizing X gives an agent A reason to φ " tell *all* agents that they should not act against the realization of X unless their reasons for doing so defeat A's reasons⁴⁸. Judgments which favor acting on aim X *intentionally* entail that all judgments involving aims which are not co-possible with X are incorrect⁴⁹. Thus, normative judgments

⁴⁸ That, at least, is what the judgment *tells* us. Whether we accept the 'grounding' judgment as correct or not is a separate issue. These entailments combined with the Humeans' do not by themselves inform a complete constructivist procedure of scrutiny, but they are needed to avoid the problems I have shown above.

⁴⁹ My proposal does not require that we think of correctness as binary. We may speak of extents of correctness in order to capture the possibility of irresolvable moral dilemmas. On my proposal, an irresolvable dilemma could be conceived as a case of two judgments with jointly unrealizable aims which are both partially correct to the same extent.

make practical reasons claims on agents even when they are not judgments about what the agent herself has reason to do⁵⁰.

4.3. How intentional entailments enable normative thought and talk

By making it the case that every judgment about other agents' reasons entail judgments about our reasons, the former can independently inform our conversational dispositions, allowing judgments about other agents' reasons to enter into the common ground. Likewise, judgments about others' reasons yield at least negative prescriptions for us as well.

To go back to Susan and Joe: if Susan were to reason in accordance with the above standards, she would judge that the value of getting her car fixed gives Joe at least a *pro tanto* reason to help. Susan's judgment that she should recruit Joe's help makes a claim on him. Her aim involves Joe's actions directly. If Joe does not help her, he is acting against her reason-giving aims – there is no middle ground between helping and hindering then. So 'not acting against the realization of her aim' involves acting in its favor. Conversely, if she judged that Joe's reasons *not* to stop defeated all the considerations offered in favor of it, the standards set by intentional entailments would bind her to also judge that she should not try to get him to stop, as that may prevent Joe from acting on those reasons. She can only justify herself in making this request of him to the extent that she judges that *he* has most normative reasons to act as she requests⁵¹.

⁵⁰ Which is not to say that what grounds facts about reasons (if they can be grounded at all) is independent of us and our attitudes. Many anti-realist theories of value are compatible with my proposal (see fn. 37).

⁵¹ These reasons are merely *pro tanto* largely because we deliberate under incomplete information – about the non-normative facts, about other agents' aims, etc. Imagine, as it is more natural, that Susan did not know the first driver coming her way. This driver – Bob – turns out to be rushing a sick child to the hospital. Susan thinks that the value of being helped by the first person to drive by gives her *pro tanto* reasons to signal to Bob that he should help her; Bob sees this, but still thinks the value of getting the sick child to the hospital as quickly as possible gives him stronger reasons to drive straight to the hospital. Her aim, 'being helped by the first person to drive by', is incompatible with Bob's aim of 'getting the sick child to the hospital as quickly as possible'. Both inform defeasible reasons: had Susan known, and assuming she would agree with Bob that getting the sick child to the hospital is more valuable than getting help from the first person to drive by, it would be most coherent of her to judge that she did not in fact have a reason to tell Bob that he should stop by signaling to him.

Intentional entailments can therefore inform agents' dispositions in much the same way instrumental entailments do. Going back to the example of the flight to Rome, the judgment that one should get to Rome immediately motivates an instrumentally coherent agent to get on a plane in light of the fact that doing so is the only way to get there quickly. Likewise, the judgment that Joe should stop motivates an intentionally coherent Susan to get him to do so – which will often involve telling him that he should stop. In this picture, the normative judgments that motivates her to tell Joe that he should stop are no longer exclusively about her own reasons. Susan acts as if she held both that she should get him to, and that he should in fact stop. So if Joe does stop by virtue of accepting what she told him about his reasons, they will have carried out joint normative inquiry: the common ground of their conversation will have been updated with the shared presupposition that Joe should stop.

When Susan tells *herself* that Joe should stop, the aim she prescribes for Joe are also informed by what she judges to be her own reasons to act, and vice-versa. Whether it be the receipt of payment by Joe, or getting her own car fixed by him, she takes *any* agent to have a *pro tanto* reason to act only on aims that are jointly realizable with them. She cannot coherently take herself (or any other agent) to have reason to prevent Joe from doing what she thinks he should. If her reasons defeated the ones in favor of Joe acting a certain way, they would also change her judgment about what Joe should do. Through intentional entailments, all normative judgments are necessarily prescriptive for all agents.

Susan can thus incorporate judgments about other agents' reasons into her deliberations in a way that preserves their prescriptive contents. She can also express these judgments as statements of fact rather than mere commands when she engages in moral discourse.

Humeans' ambition to get standards of correctness purely out of instrumental practical reason is motivated by their skepticism that there are any viable alternatives which do not rely on a realist/externalist

picture of moral facts, which Humeans have their own reasons to reject⁵². I cannot fully address these concerns here. But I believe that they overestimate the challenges of incorporating other kinds of constitutive entailments which are not purely logical or instrumental, and underestimate the challenges they themselves face for limiting their constructivist procedure this way. These limitations are what lead them to validate the conjunction of judgments involving conflicting aims. Admittedly, a view which includes my proposed entailments no longer immediately tells us how to adjudicate moral disagreement between agents with different values. But that alone does not mean my proposal cannot be supplemented by accounts of how we might⁵³.

5. Conclusion

Humean constructivism in meta-ethics does not take aims into account as much as, or in the way that it should. One particularly problematic commitment of the Humean view is that it allows for multiple normative judgments to be correct even when they prescribe acting on incompatible aims. (By incompatible, I mean that agents cannot jointly realize those aims no matter what they do – including if they did what correct judgment tells them they should.)

To show that this commitment represents a problem for the Humean, I took an indirect route. Instead of pumping intuitions against that upshot and in favor of my way of seeing things, I challenged what I see as the source of the problem in the Humean view: the fact that their correctness-constituting entailments only relate judgments that involve the same agents. This locks agents who reason correctly according to the Humean view into a kind of deliberative narcissism. They are only able to talk and think normatively about other people's reasons through normative judgments they make about themselves.

⁵² See Street's debunking argument in her well-known *Darwinian dilemma* (2006), or Bernard Williams' case for internalism (1981; 1995).

⁵³ Korsgaard (2009: 206), Bratman (2009: 417ff), and even expressivists like Gibbard (1992: 22) reference analogues of intentional coherence as standards for practical reasoning in their writing.

Against the letter of Humeans' own picture of normative judgment, moral discourse becomes purely prescriptive, as Humean agents are largely incapable of adding normative propositions to the conversational common ground. Thinking about other agents' reasons in turn becomes purely descriptive, as judgments about others carry no normative weight.

It is much more plausible to think that correct practical reasoning does not force this kind of self-centeredness on us. A lot about our ability to deliberate about other agents' reasons in *sui generis* normative terms depends on our ability to adjudicate the general validity of those reasons. When we make normative claims about other agents, we make claims about whether the prescribed aim should be realized in the world we all share. This allows us to move away from the prescriptivism of Humean conversation, treating the claims we make on others as part of a joint inquiry into the normative facts. In turn, thinking about other agents' reasons can be just as normatively forceful as judgments we make about our own reasons. This is because, according to my proposal, judgments about other agents' reasons and ours can entail each other.

Lastly, the case for my proposal makes use of a similar constitutivist strategy, but unlike Humean constructivism, it allows us to make sense of how *all* of our normative judgments might cohere with one another – rather than thinking of rational coherence as something that applies only to judgments about what we ourselves have reason to do. The proposal vindicates the connection between what I believe to be the two essential functions of practical reasoning: thinking about what we ourselves should do, and making claims on others. Judging oneself to have reasons to interact with someone else constitutively involves making claims about their reasons.

Works cited

- Bedke, Matthew S. "Might all normativity be queer?" *Australasian Journal of Philosophy* 88, no. 1 (2010): 41-58.
- Blackburn, Simon. *Spreading the word: groundings in the philosophy of language*. Oxford University Press, 1984.
- Blackburn, Simon. "Hume on the mezzanine level." *Hume Studies* 19, no. 2 (1993): 273-288.
- Bratman, Michael E. "Intention, practical rationality, and self-governance." *Ethics* 119, no. 3 (2009): 411-443.
- Dorsey, Dale. "A Puzzle for Constructivism and How to Solve It." *Constructivism in practical philosophy* (2012): 99.
- Gibbard, Allan. *Wise choices, apt feelings: A theory of normative judgment*. Harvard University Press, 1992.
- Gibbard, Allan. "Morality as consistency in living: Korsgaard's Kantian lectures." (1999): 140-164.
- Korsgaard, Christine M., *The sources of normativity*. Cambridge University Press, 1996.
- Korsgaard, Christine M. *Self-constitution: Agency, identity, and integrity*. Oxford University Press, 2009.
- Korsgaard, Christine M. "Valuing our Humanity" *Respect*, Oxford University Press, forthcoming.
- LeBar, Mark. "Aristotelian constructivism." *Social Philosophy and Policy* 25, no. 1 (2008): 182-213.
- Mackie, John Leslie. *Hume's moral theory*. Routledge, 2003.
- Nagel, Thomas. *The view from nowhere*. Oxford University Press, 1989.
- O'Neill, Onora. *Constructions of reason: Explorations of Kant's practical philosophy*. Cambridge University Press, 1989.
- Scanlon, Thomas. *What we owe to each other*. Harvard University Press, 1998.
- Scanlon, Thomas M. *Being realistic about reasons*. Oxford University Press, 2014.
- Stalnaker, Robert. "Common ground." *Linguistics and philosophy* 25, no. 5-6 (2002): 701-721.
- Street, Sharon. "A Darwinian dilemma for realist theories of value." *Philosophical Studies* 127.1 (2006): 109-166.
- Street, Sharon. "Constructivism about reasons." *Oxford studies in metaethics* 3 (2008): 207-45. Online copy.
- Street, Sharon. "In defense of future Tuesday indifference: Ideally coherent eccentrics and the contingency of what matters." *Philosophical Issues* 19, no. 1 (2009): 273-298.
- Street, Sharon. "What is constructivism in ethics and metaethics?" *Philosophy Compass* 5.5 (2010): 363-384.

Street, Sharon. "Objectivity and truth: You'd better rethink it." *Oxford studies in Metaethics* 11 (2016): 293-334.

Streumer, Bart. *Unbelievable Errors: An Error Theory about All Normative Judgements*. Oxford University Press, 2017.

Velleman, J. David. *How we get along*. Cambridge University Press, 2009.

Williams, Bernard. "Internal and External Reasons." *Moral Luck* (1981): 101-113.

Williams, Bernard. "Internal Reasons and the Obscurity of Blame" *Making sense of humanity and other philosophical papers, 1982-1993*. Cambridge England New York: Cambridge University Press (1995): 35-45