

Communication for Expressivists*

Alejandro Pérez Carballo and Paolo Santorio

1 Introduction

According to metaethical expressivists, moral statements do not aim at describing the world. The function of moral talk is instead to express nonrepresentational mental states. As is well-known, expressivists face the challenge of giving a semantics for moral language that is compatible with these claims—and the jury is still out on whether this challenge can be met. In this paper, we focus on a related, but different question: how can the expressivist account for the assertion and the uptake of normative claims?

Suppose that Anne and Brad are discussing the permissibility of tax evasion. Anne utters ‘Tax evasion is wrong.’ She does this because she has an intention to produce in Brad a certain attitude about tax evasion. Upon hearing Anne’s utterance, Brad recognizes her intention and acts accordingly: he either adopts the relevant attitude or challenges Anne’s claim. All this has the semblance of a rational and purposeful activity. Can the expressivist account for this intuition?

To see that the expressivist has a *prima facie* problem, consider how an analogous question is answered for the case of uncontroversially descriptive discourse. Suppose that Anne asserts that a common acquaintance, Tom, evaded taxes last year, and Brad comes to believe the corresponding proposition. How can we explain Anne’s assertion and Brad’s uptake of the relevant information? Answers may vary, but they will share a common core. Anne makes her assertion because she thinks this is a good way of imparting some relevant knowledge to Brad. Brad comes to believe what Anne says because it promotes his goal of learning information about the world. This simple explanation crucially appeals to notions like truth, knowledge, and the actual world. But it is part of the expressivist’s picture that, at least in an important sense, normative statements are not truth-apt and don’t convey information about the world.

Some expressivists suggest that we engage in conversation about normative matters to achieve a kind of coordination. This, they suggest, is the goal of conversation about both normative and non-normative matters: to get participants’ attitudes to align, whether those attitudes are beliefs or attitudes of another kind. But even if this is right, it is far from a complete account. If coordination was all we wanted out of conversation,

*Forthcoming in *Ethics*. Please consult the published version before quoting.

our communicative practices would look very different. Flipping a coin is a perfectly acceptable way to coordinate. But we don't solve normative disagreements via coin flips or similar randomization strategies. Hence speakers are not *merely* looking to coordinate when engaging in normative talk. Something more than mere coordination must be at stake. For the case of communication about non-normative matters, truth can be again invoked to the rescue. Speakers want to coordinate by having their beliefs align with what things are actually like. But, again, the expressivist cannot appeal to truth in her account of normative talk.

In this paper, we provide the beginning of an account of communication on behalf of the expressivist. We start from basic assumptions about the rationality of speakers and about assertion. These assumptions, we argue, entail that speakers engaging in normative communication must assume that their claims are subject to a standard of correctness, similarly to descriptive claims. More precisely, this is our main contention:

In any conversation where certain minimal assumptions are satisfied, it is presupposed that there is a (unique) normative standard on which the participants' attitudes ought to converge.

The *ought* here is a practical one. It concerns the rational requirements that apply to agents, in view of their goal of taking part in communication. Our claim, then, is this: to successfully engage in communication, speakers have to presuppose that they're subject to a requirement to converge on a joint standard that applies to them in virtue of their being communicators. We explain the precise content of this requirement, together with the 'minimal assumptions', the notion of a presupposition, and the nature of the attitudes at stake, in the course of our discussion.

We should point out that we don't aim at giving a full account of communication on behalf of the expressivist. We take for granted certain basic facts about assertion that an expressivist will need to explain further.

Throughout the paper, we exploit a specific formal framework for modeling communication, namely the one proposed by Robert Stalnaker and refined by many others. Our argument does not depend on the specific formal features of Stalnaker's framework. We use this framework because it makes our argument clearer and more precise. But ultimately our argument could be run independently of the overall framework. All that we strictly need are two assumptions. The first is that to make an assertion is to propose an update of certain belief-like attitudes of the speakers. Hence, on this picture, communication is first and foremost a way of exchanging information (broadly construed, to include also normative information). The second assumption is that the relevant belief-like attitude should be understood in terms of Stalnaker's notion of acceptance, which we discuss below. We think these assumptions are essentially correct: they capture what happens in the great majority of ordinary conversations. (In fact, we think that the whole of Stalnaker's framework provides a useful model of conversations

between real agents—though, like every model, it involves some level of idealization.) Hence, while our argument is framed in formal terms and the notions of assertion, communication, and conversation we use are technical, our conclusion is intended to apply to actual normative conversations between real agents.

2 Background: expressivism and communication

Throughout the paper, we take expressivism about normative discourse as our running example. Ultimately, we would want our conclusions to apply to expressivism about a large variety of discourses. But different varieties of expressivism call for tweaks and extensions of our claims. Hence, for the purposes of this paper, we set them aside.¹

2.1 Expressivism

We start from a minimal conception of expressivism about normative discourse, consisting of a negative and a positive claim:²

Minimal Expressivism

- a. normative claims are not apt for describing, stating, or reporting facts;
- b. normative claims express a non-cognitive (nonrepresentational) attitude of some sort.

Of course, to get a proper theory out of it, this minimal characterization needs to be fleshed out. One needs to say more about, first, the nature of the non-cognitive attitudes in play and their role in a general philosophy of mind; and second, the expressivist's semantics for normative discourse.

For present purposes, we will work with the version of expressivism developed by Allan Gibbard in his (1990). Gibbard's expressivism has a number of virtues. It is a general theory of normative discourse; it yields a semantics that is compositional and fully compatible with standard syntactic views; it yields a simple account of logical consequence for normative discourse.³ Moreover, Gibbard's semantics for normative dis-

¹In particular, so-called expressivism about epistemic and probabilistic discourse (see, among others, Yalcin, 2011, Swanson, 2011, Rothschild, 2012, Yalcin, 2012) might require some substantial changes to the argument. It's not clear to us whether, on these accounts, the update of shared information happens via set-theoretic intersection of some formal objects, as it happens for Gibbard, or whether they employ a different update operation (as it happens in the account given by Veltman, 1996, which is an obvious source of inspiration for these views). Hence it's not clear to us whether these accounts can be modeled as a generalization of the Stalnakerian model of assertion we adopt.

²For a similar characterization, see Price, 2011. Notice that we will reserve the term 'expressivism' to what Price calls '20th century expressivism'.

³In particular, we take Gibbard's semantics to have the best chances of overcome the notorious Frege-Geach problem. For the initial formulation of the problem, see Geach, 1960, 1965. For an overview of the

course dovetails well with the framework for modeling communication we will use, namely the one developed by Stalnaker, 1978.

2.2 Gibbard's norm-expressivism

Gibbard's account is formulated in terms of an attitude which he calls 'accepting a norm': this is the kind of attitude that determines what an individual regards as mandated, permissible, or forbidden. Your judging that cannibalism is wrong amounts to your accepting a norm that forbids cannibalism; Tom's judging that tax evasion is okay amounts to his accepting a norm that allows for tax evasion. Gibbard doesn't define the notion of accepting a norm: rather, he assumes that this attitude will be part of a (yet to come) fully developed empirical psychology, on a par with beliefs and desires. The notion of accepting a norm is supposed to play a central role in an evolutionary explanation of individuals' coordination in a social context.

One crucial part of Gibbard's enterprise consists in specifying a formal model for the contents of normative attitudes. To do this, Gibbard employs *complete systems of norms*. Think of a complete system of norms n simply as (something that determines) a three-way partition of possible courses of action: those that are forbidden according to n , those that are permitted but not required according to n , and those that are required according to n . Complete systems of norms (henceforth, for short: norms) play the same role in characterizing normative attitudes that possible worlds play for non-normative attitudes in possible worlds semantics. If you believe that Tom is a cannibal, we characterize the content of your belief state with worlds where (among other things) Tom eats human flesh. Similarly, if you accept that cannibalism is wrong, Gibbard models the content of your normative attitudes with a set of norms all of which forbid (among other things) eating human flesh.

With this model of mental content in place, it is easy to formulate a semantics for normative language that works in the usual recursive way.⁴ So far, we have assumed that descriptive attitudes make distinctions between possible worlds and that normative attitudes make distinctions between norms. But it is best to let all sentences in a language denote formal entities of the same kind, for a number of reasons (for example, accounting for 'mixed' sentences like *If Tom eats people, he does something wrong*). So instead of assigning semantic values of different kinds to each fragment of the language, we assign *sets of world-norm pairs* to each sentence. Hence the semantic values of

- (1) Tom is a cannibal.

literature and discussion, see Schroeder, 2008a,b, 2010, among many others. As we note below, there is still considerable skepticism that Gibbard's picture can provide a real solution to the problem.

⁴Gibbard's own way of doing so yields the wrong results for attitude reports. But those are wrinkles that can be easily ironed out. Cf. Dreier, 1999, p. 571.

(2) Eating people is okay.

are given by, respectively:

(3) $\{\langle w, n \rangle: \text{Tom is a cannibal in } w\}$

(4) $\{\langle w, n \rangle: \text{Eating people is permitted by } n\}$

That said, we will proceed under the simplifying assumption that all purely normative sentences get assigned a set of norms as semantic value. We do this just to reduce clutter. But it's easy to recover sets of world-norm pairs from our semantic values: simply identify each set of norms A with the set of world-norm pairs $\langle w, n \rangle$ such that n is in A .

Two final clarifications are in order. First, recent literature has raised some foundational questions about Gibbard's framework. The main worry is that Gibbard cannot specify (in a non-stipulative and non-circular way) what it is for two norms to be inconsistent. Here, we set this worry aside—we assume that the problem can be solved one way or another.⁵ Second, for the purposes of this paper, we assume that formal objects like (3) and (4) work both as the semantic values recursively assigned to (1) and (2) and as the contents of utterances of (1) and (2). There are good reasons to think that this is a conflation,⁶ but one that is harmless for our purposes.

2.3 Assertion, communication, common ground

Throughout the paper, we rely on a broadly Stalnakerian picture of communication (see e.g. Stalnaker, 1973, 1978, 2002, 2014, as well as Lewis, 1979b). On this picture, conversation takes place against a background body of information—the 'common ground'. Roughly, the common ground represents what is commonly taken for granted for the purposes of the conversation. The purpose of an assertion is to expand this shared stock of information. A successful utterance brings about that the proposition uttered gets added to it.

We represent a body of information as a set of possible worlds—the worlds that are compatible with that body of information. Following Stalnaker, we call the set of worlds modeling the background information in a particular context the 'context set'. We also think of the content of a declarative utterance as a set of possible worlds—those worlds in which the utterance is true. Finally, we model the effects of an assertion on the context set as set-theoretic intersection: The effect on the context set of an utterance whose content is modeled with a set S of possible worlds is that of eliminating from the context set those worlds not in S .

⁵Or perhaps, with one of the authors, that the problem turns out to be spurious on closer analysis. See Pérez Carballo, 2012.

⁶For discussion of this point, see Dummett, 1981; Lewis, 1980; Ninan, 2010; Rabern, 2012.

Consider a simple example. Zoe and Yael are having a conversation about Tom. They jointly assume a number of propositions about him: for example, the proposition that Tom lives in Boston. We model the context set of their conversation as a set of worlds in which (among other things) Tom lived in Boston. Now, suppose Yael utters the sentence ‘Tom didn’t pay taxes last year’. If her assertion is successful, the context set will now get updated with this new piece of information: hence the new context set contains only worlds where Tom lives in Boston and didn’t pay taxes last year.

We can specify what it is for some information to be “commonly taken for granted” in more detail. First, we need to introduce the notion that Stalnaker (2002) dubs ‘acceptance’. This notion is entirely distinct from Gibbard’s notion of accepting a norm; the terminological similarity is just an unfortunate coincidence (we worry about how to deal with this below). Acceptance is, quite simply, the attitude of taking a proposition for granted for the purposes of a conversation.

In at least some cases, acceptance diverges in interesting ways from belief. Consider this example:

Smith and Jones find themselves in a conversation. They both believe (suppose, truly) that, many years ago, Smith tried to steal money from Jones in a moment of desperation; moreover, they both believe that they believe this, and believe that they believe that, etc. In short, that proposition is the object of common belief between them. However, they both prefer avoiding to acknowledge that proposition. So, for the purposes of the conversation, they behave as if they did not believe this proposition.

That Smith once tried to steal money from Jones is not accepted and is not validated by the context set of their conversation. Notice that acceptance has a peculiar ‘transparent’ nature: if an agent represents herself as accepting p , then, as a matter of fact, she does accept p for the purposes of the conversation. One cannot lie about one’s acceptances. The transparent nature of acceptance will play a crucial role for us.

The contents of the common ground are defined on the basis of acceptance. In particular, the common ground will be identical with what is commonly believed to be accepted in a conversation. Common belief is an iterated notion of belief: it is common belief that p just in case all the members of the conversation believe that p , all believe that they believe that p , etc. So, in summary:⁷

⁷Here we go along with the notion of common ground from Stalnaker, 2002:

The more general notion of common ground should not be just an iterated version of a broader notion of acceptance. Successful communication is compatible with presuppositions that are recognized to be false, but the information that they are being presupposed must be actually available, and not just assumed or pretended to be available. Even the liar, if he really intends to communicate, has to believe that the information needed to interpret his lies will really be common ground. (p. 716)

In recent work Stalnaker reverts to a picture of common ground that uses just iterated acceptance—see Stalnaker, 2014. Going along with this alternative picture would make no difference for our purposes.

It is *common ground* that p in a group if all members accept (for the purpose of the conversation) that p , and all believe that all accept that p , and all believe that all believe that all accept that p , etc.

The context set is just the set of worlds that validate all and only the propositions that are common ground.

Using similar tools, we can also define a notion of speaker presupposition. Differently from common belief, presupposition is an individual attitude. Each speaker enters a conversation with their own presuppositions. Roughly, a presupposition is a proposition that the speaker takes for granted in the course of the conversation. We can capture this pretty simply with the current apparatus by identifying the presupposition that p with the belief that it is common ground in the conversation that p :

Speaker S *presupposes* that p in a conversation iff S believes that it is common ground among participants in the conversation that p .

For the purposes of our argument, we often take presupposition, much more simply, as mutual acceptance. But we revert to the more precise notion where it matters.

Gibbard-style expressivists can take on board this general picture of communication. There is a terminological complication, however. The notions of acceptance, common ground, and presupposition have been introduced in a truth-conditional framework. Hence these notions concern descriptive contents. We would need to introduce new terminology to capture counterparts of them that also apply to normative contents. But this would involve a lot of clutter, so we'll take a shortcut. We are going to understand the standard vocabulary of theories of assertion and communication in a broader, nonstandard way. We take a proposition to be modeled either by a set of worlds or by a set of norms.⁸ We use 'belief' to pick out both descriptive beliefs and attitudes of accepting a norm in Gibbard's sense. Hence we'll talk about a subject believing that (say) tax evasion is wrong, while still understanding this attitude in a noncognitivist way. Similarly for acceptance and presupposition.⁹

We can now summarize our Stalnakerian picture of conversation, appropriately enlarged to encompass also Gibbardian contents. In conversation, speakers try to influence what others accept—where *to accept that p* is to act as if one believes that p . The contents of the speakers' mental states can now be modeled via two parameters, i.e. worlds and norms. Hence the possibilities that are part of the common ground are different from those we find in the descriptivist model. But all the rest of the apparatus remains the same. In particular, we still model the common ground as (roughly) the set of possibilities (norms) that are left open by speakers' attitudes, and we still model the effect of an assertion on the context set by intersecting it with the semantic value of the sentence asserted.

⁸In full generality: a set of world-norm pairs.

⁹For discussion of this and related terminological issues, see Schroeder, 2010, p. 85ff.

Before proceeding, let us address an objection. One might worry that the common ground model is incompatible with Gibbard's norm expressivism.¹⁰ The notion of assertion, the objection goes, is closely connected to notions like truth and knowledge, which belong in a factualist picture of language.¹¹ Perhaps there are no formal obstacles to extending Stalnaker's formalism to Gibbardian contents. But doing so involves dropping some of the main philosophical commitments of expressivism.

We grant that there is a construal of assertion on which it bears substantial, perhaps conceptually necessary connections to truth and knowledge. But here we rely on a 'thinner' construal of assertion. We understand assertion as the speech act whose functional role is to update the common ground in a certain way—in particular, assertion is the speech act that updates the common ground by intersection. Given our expressivist-friendly understanding of acceptance, and the plausible claim that utterances of declarative normative sentences can update speakers' acceptances, the expressivist is entitled to characterize utterances of declarative normative sentences as assertions in this sense.

Notice that this explication of assertion makes no appeal to notions like truth and knowledge. Instead, the thin construal we use ties assertion to linguistic form. In general, assertions are the speech acts that are performed by declarative sentences. We say 'in general' because not all utterances of declarative sentences will do: sometimes we can utter a declarative sentence with no intention to produce a speech act (for example, just to practice our pronunciation). Assertions are the speech acts that are performed via those utterances of declarative sentences that are intended to update the common ground.

2.4 Issues

In any conversation, there are a number of questions that are of mutual interest to the participants in the conversation. These questions can be used to define a notion of *relevance*: a proposition is relevant in a context *c* iff it does not say more than what is needed to answer all of the questions. And, more importantly, they can be used to provide a less idealized characterization of the common ground.¹²

If we model the possibilities compatible with the common ground with a set of worlds (the context set), we can model each question of interest—or 'question under discussion'—as a *partition* of this set: a collection of subsets of the context set that are pairwise disjoint and jointly exhaustive. Two worlds are in the same set iff they agree on the complete answer to the question.¹³


¹⁰Thanks to an anonymous referee for pushing us to consider this and related issues.

¹¹For some classical accounts vindicating these connections, see Dummett, 1959 and Williamson, 1996, 2000.

¹²Cf. Hulstijn, 1997 and Roberts, 2012.

¹³Cf. Groenendijk and Stokhof, 1984; Hamblin, 1958, 1973; Lewis, 1988, *inter alia*.

For illustration, suppose that Yael and Zoe are having a conversation with three questions under discussion: whether Tom paid taxes last year; whether he concealed anything from the IRS; and whether he is going to be audited. Here is a graphic representation of the partitions corresponding to each of these questions:



Given a set of questions under discussion, we can define a unique partition, which groups together worlds that agree on the answer to *all* the questions under discussion.¹⁴ In the figure above, the partition is composed of those regions that are not divided by any of the lines.

On this picture, the goal of conversational participants is no longer to locate the actual world within logical space, but (more modestly) to individuate the true answers to all of the questions under discussion. This is equivalent to ruling out all the cells in the partition, aside from the one that contains the actual world. For example, suppose that Tom has indeed paid taxes, did not conceal anything from the IRS, but will nonetheless be audited. Then Yael and Zoe’s goals are accomplished when all cells other than the bottom left L-shaped cell are ruled out.¹⁵

Once we track questions under discussion, the possibilities that make up the common ground can be thought of not as possible worlds, but as cells of the partition induced by the questions under discussion. Following the terminology of (Savage, 1972),¹⁶ we can call these cells ‘small worlds’ (they are called ‘small’ because they contain strictly less information than full-blown possible worlds.) Small worlds are different from possible worlds, but the common ground model works exactly in the same way, whether we use the ones or the others. In both cases, assertion works by ruling out open possibilities.

In what follows, we take small worlds to be the basic units of the context set. What small worlds are in play is a function of context: we assume that, for each context, there

¹⁴More precisely: this is the coarsest partition that is a refinement of all the questions under discussion.

¹⁵This modification of the common ground framework is very much in the spirit of the original proposal. As Stalnaker puts it (1981, p. 136): “The alternative possibilities used to define propositions must be exclusive alternatives which are maximally specific, relative to the distinctions that might be made in the context at hand. [...] One might think of possible worlds as something like the elements of a partition of a space, rather than as the points of the space. The space might be partitioned differently in different contexts, and there might be no maximally fine partition.”

¹⁶See also Joyce, 1999 for recent discussion.

are a number of questions under discussion that induce a partition of the context set into small worlds. Further questions under discussion might be added in the course of the conversation, both explicitly, i.e. by uttering an interrogative sentence, or implicitly.¹⁷ But whether an issue becomes an question under discussion ultimately depends, on our view, on whether all speakers intend to settle that question in conversation:

The set of questions under discussion in a conversation is the set of questions Q such that, for each Q , all speakers intend to engage in inquiry to settle whether Q .

In short: the live questions are those that speakers intend to discuss.¹⁸ As we said above, the questions under discussion in a conversation generate the small worlds in the context set. Hence, for each small world in the conversation, all speakers involved are willing to take up the question whether that particular small world in the conversation should be ruled out.

All considerations that we've just made can be reproduced, *mutatis mutandis*, for Gibbardian norms. We assume that the context set in conversations about normative matters consists of 'small norms'—sets of complete systems of norms that are grouped according to how to settle all normative questions under discussion.

Even though we take small worlds as the basic units of the context set, to avoid clutter we keep using the simple terms 'world' and 'possibility'. Most of what we say holds both for small worlds and full-blown possible worlds. But we will remind you of the explicit assumption that small worlds are the ones in play in places where it matters. Similarly, we will just talk about norms. As a result, when we say that conversation about normative matters involves distinguishing among complete systems of norms, we should be understood to mean 'complete only to the extent required by the purposes of the conversation'. We will remind you of the explicit assumption that small norms are the ones in play in places where it matters.

¹⁷Note that implicitly raising an issue need not automatically result in a modification of the set of questions under discussion. Suppose you have a conservative uncle who often makes in-passing claims that are off-topic, and which presuppose some of his views about (say) abortion. In conversation about different topics, may want to avoid pointless confrontation and you choose not to pick up on these remarks. In these cases, the question whether abortion is permissible has no effect on the common ground. Thanks here to an anonymous referee, to whom we also owe the example of the reactionary uncle.

¹⁸This characterization has a variant worth considering:

The set of questions under discussion in a conversation is the set of questions Q such that, for each Q , all speakers intend to engage in inquiry to settle whether Q , and all speakers believe that they intend to engage in inquiry to settle whether Q , and they believe that they believe it, etc.

The difference between these two characterizations is irrelevant for our purposes, so we stick to the simpler one.

3 The problem

We established that expressivists can borrow from descriptivists a formal framework for modeling content, assertion, and communication. But the formal framework leaves one important question unaddressed. Expressivists should be able to *make sense* of the practice of conversation. They need to explain why it's reasonable for a speaker to have certain expectations about the uptake of what she says, and why it's reasonable for a hearer to actually take up what the speaker says.

In the descriptivist version of our model of communication, the explanation goes as follows. Communication is a self-locating enterprise. Speakers aim to find out what portion of logical space they occupy. Exchanging information is just a way of jointly narrowing down the set of candidates for the actual world. But the expressivist can't help herself to a normative counterpart of the actual world. It is a basic assumption of her view that there is no such thing. Hence the expressivist owes us a story about the point of engaging in communication about normative matters, one that she can't simply extract from the model of communication she's borrowing.

Expressivists have gestured towards a way of making sense of communication that generalizes to descriptive and normative discourse alike. The idea is to think of communication primarily as an exercise in coordination. We are social creatures: we engage in conversation because coordinating on some particular family of attitudes—beliefs, say, or norm-acceptances—is likely to help us meet our goals. Here is Gibbard explicitly making this point:¹⁹

The biological function of the mechanisms underlying our normative capacities is to coordinate. Hence the psychic mechanisms that produce normative judgments are not systems of natural representations, they are coordinating systems. Their biological function is not to put something in the head in correspondence with their subject matter; it is to coordinate what is in one person's head with what is in another's. (Gibbard, 1990, p. 110)

To a first approximation, this seems like a good strategy. The point of engaging in conversation about non-normative matters is to coordinate on a body of beliefs. The point of engaging in conversation about normative matters is to coordinate on a system of norms. We work out, as a community, what to think about the world. We work out, as a community, what system of norms to accept. Communication can be understood as a way to foster such coordination.

Unfortunately, this is far from a full account. There are many ways of achieving coordination that differ from communication in crucial ways. A newly established community might coordinate on what side of the road to drive on via a random procedure

¹⁹Seth Yalcin, who endorses an expressivist account of epistemic modal talk, makes a proposal along similar lines. Cf. Yalcin, 2011, p. 310.

like flipping a coin. This is a perfectly acceptable way of achieving coordination. But we don't rely on coin flips when working out what to do, no more so than we rely on coin flips when working out what the world is like. So saying that communication is a kind of coordination is just the beginning of an account. The expressivist stills owes an account of exactly what kind of coordination this is and what constraints apply to it.

One natural thought is that, when picking a side of the road to drive on, all we want is that we all drive on the same side of the road. None of us cares *which* side that is. Perhaps the case of communication is different because agents approach communication with a number of initial beliefs (normative or not). Besides jointly coordinating on a sector of logical and normative space, they care about making it the case that this sector not be disjoint from the sector that is already individuated by their initial attitudes. In other terms, agents care about maximizing the extent to which their initial views are preserved after communication.

There is something right about this suggestion, but it also falls far short of an account. If communication functioned this way, it would be mostly an exercise in bargaining. In other words, communicating would be about proposing compromises that have a high enough payoff for all agents involved. Of course, we can and do engage in bargaining of this sort. Suppose that you and a friend care about spending the evening together, but have conflicting preferences. In that situation, it would not do to just flip a coin to settle each detail of your plan. There would be some give and take: you eat at the diner rather than the fancy restaurant, as per your own preferences; but you go to the opera rather than the roller derby, as per your friend's. But this is not what we do when we discuss descriptive and normative matters. Then we try to reach a set of shared attitudes by giving reasons to each other. The idea that we could compromise by just trading up some of our beliefs and acceptances for agreement is out of place.²⁰

Thus we grant that coordination should play a part in a general account of communication. But a desire for coordination on our beliefs or our normative acceptances cannot explain the complexities of our communicative practices. The expressivist needs more to get a full account.

4 Explaining communication

We now try to spell out what extra element needs to be in place in normative conversations, besides an intention to coordinate. We argue that this extra element can be extracted from some general assumptions that speakers need to make in order to engage in conversation.

²⁰Incidentally, this observation also shows that the expressivist's acceptances, although they are a conative attitude of some sort, cannot be construed too closely on the model of preferences. We take it that this is also one of the main lessons of Dorr, 2002—though, unlike Dorr, we don't take this disanalogy to doom expressivism.

We start by examining the assumptions agents need to make in order to purposefully and rationally engage in communication about normative matters. We argue that, to do this, agents must take for granted that there is a kind of objectivity about normative matters. As a result, in each context, speakers presuppose that they are required to converge on a unique norm, at least insofar as they choose to remain engaged in communication. This will explain why our practice of communicating and debating normative claims is analogous to our practice of communicating and debating non-normative claims. In normal conversations, speakers presuppose that there is an actual world: hence they presuppose that there is a point in the context set to which, at least ideally, their beliefs about the world should converge. Similarly, we claim that they jointly presuppose that there is an ‘actual’ norm—a norm to which, at least ideally, their beliefs about normative matters ought to converge.

4.1 The claim: convergence

Let us start by introducing a bit of formalism. We use ‘ n ’ and ‘ s ’ as variables ranging over systems of norms and speakers, respectively. (More precisely, ‘ n ’ should be understood as ranging over systems of norms that are complete *only to the extent required by the purposes of the conversation*.) We use ‘ Acc_s ’ to denote the set of systems of norms compatible with what s accepts, and we use ‘ Acc_s ’ as an ‘acceptance’ operator, saying that the embedded proposition is accepted by s . We can formulate our main conclusion as follows:²¹

CONVERGENCE. The following is common ground in any conversation:

$$(\exists n)(\forall s)\text{Ought}(\{n\} = Acc_s).$$

We argue that CONVERGENCE follows—surprisingly, perhaps—from some minimal assumptions about the rationality of agents engaging in communication and our practice of assertion.²²

Let us say a bit more about the nature of the *ought* figuring in CONVERGENCE. This *ought* is supposed to capture norms regulating communication between rational agents: i.e., norms that specify what a rational agent who has the goal of engaging in a communicative exchange ought to do. Hence it is an *ought* of practical, as opposed to epistemic,

²¹Henceforth, we drop the qualifier ‘normal’ for the sake of readability.

²²Notice that CONVERGENCE does not amount to a capitulation to a form of realism. A realist will presumably think that there is a unique norm (the ‘true’ norm) to which, in any ‘normal’ conversation, participants’ acceptances ought to converge. She may even think that there is a unique norm such that, in any ‘normal’ conversation, it is common ground that the participants’ acceptances ought to converge to that norm. But neither of these claims is entailed by CONVERGENCE.

normativity.²³ A rough gloss of ‘ A ought to φ ’ would go like this: in view of the requirements of practical rationality applying to her qua agent involved in a communicative exchange, A ought to φ . So far as we can see, the normative force attaching to this *ought* is the same as that of Grice’s Cooperative Principle:

Make your conversational contribution such as is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which you are engaged.

Grice claims that anyone who cares about the central goal of communication is rationally required, in light of their own goals as a communicator, to make their conversational contribution adequate in certain ways. Similarly, we claim that, in a normal conversation, speakers presuppose that there is a unique norm such that they are rationally required, in light of their own goals as communicators, to have their attitudes converge on that norm.

Before moving on, we should emphasize again that ‘norm’ in CONVERGENCE is short for ‘small norm’, i.e. a set of norms that is a cell of the partition generated by questions under discussion. Hence CONVERGENCE amounts to the requirement that, in ‘normal’ conversations, speakers presuppose that there is a unique answer to all questions under discussion such that that they ought to converge on that answer. This falls short of a presupposition that there is a fully specified normative standard that agents ought to converge on, but it still is a substantial step towards endorsing a kind of presupposition of objectivity in normative conversations. Beyond being a step in the direction of objectivity, this arguably establishes a full parallel with conversations about non-normative matters. Also in that case, speakers presuppose that there is a unique ‘small world’ (the one containing the actual world) to which their acceptances ought to converge. Plausibly, in real conversations they do not presuppose that there is a unique, maximally specific, possible world to which their acceptances ought to converge. Such a presupposition would presumably be in conflict with Gricean assumptions about relevance. Questions under discussion are all that is relevant in a conversation, and speakers have no interest in irrelevant information.²⁴

4.2 Not anything goes

Our argument for CONVERGENCE begins by establishing a weaker claim. In intuitive terms, this claim is that, in normative conversation, it is common ground that there

²³There are traditional worries about the idea that practical normativity may apply to an agent’s cognitive states. These worries are connected to skepticism about the possibility of doxastic control. But these worries should not apply to a notion of acceptance (see below). Acceptances are propositions that agents choose to take for granted for the purposes of the conversation. Agents may not be able to decide what to believe, but they’re able to decide what to accept in a conversation.

²⁴Thanks to an anonymous referee for pushing us to consider the apparent conflict between CONVERGENCE and the maxim of relevance.

is the possibility of a mistake—i.e., it is common ground that, by participating in normative conversations and ruling out normative possibilities in accordance with what is said, speakers might rule out a normative possibility they ought not rule out. This claim can be formalized as follows:

CAN GO WRONG. The following is common ground in any conversation:

$$(\exists n)(\forall s)\text{Ought}(n \in \text{Acc}_s).$$

This premise can be regarded as a minimal requirement of objectivity: it must be common ground that at least one of the normative possibilities that speakers regard as open ought not be ruled out.

Notice that this is much weaker than the claim that speakers’ attitudes should converge on a unique norm. It only requires that they keep regarding one of the possibilities as open: i.e., that they not be allowed to accept normative propositions incompatible with it.

Our argument for CAN GO WRONG is somewhat involved. The basic point, however, can be stated quickly and informally. Consider a conversation between two agents, *A* and *B*, and suppose that the following is the case:

The ‘anything goes’ situation. *A* accepts that there is no norm such that *A* ought not rule out that norm. In other words, the following is the case, according to *A*:

$$\neg\exists n\text{Ought}(n \in \text{Acc}_A)$$

We claim that, in the ‘anything goes’ situation, it would be pointless for *B* to engage in conversation with *A*.

It helps to consider an example. Suppose *A* and *B* have narrowed down the common ground to a set of two norms (by which, once more, we mean two cells in the partition determined by their questions under discussion). On one of these norms cannibalism is wrong, while on the other it is okay. And now, suppose that *A* and *B* are indeed in an ‘anything goes’ situation, so that, according to *A*, she is allowed to accept that cannibalism is wrong or to accept that cannibalism is not wrong. *B* might then say (for example) “Cannibalism is wrong”. *A* can agree that there’s nothing wrong at all with *B*’s saying or thinking that—she would not be willing to say that *B* said anything she ought not have said. Further, *A* could agree that there is no reason for *A* not to accept *B*’s answer—she takes herself to be allowed to rule out the norm according to which cannibalism is okay. And yet, she may stick to her belief that cannibalism is okay, since she also takes herself to be allowed to rule out the norm according to which cannibalism is wrong.

We claim that, in this kind of situation, it would make no sense for *B* to engage in conversation with *A*. (Or at least: it would make no sense for *B* to engage with *A* in a conversation as long as the question whether cannibalism is wrong is a question under discussion. We get back to this shortly.) There would be no point to it. Hence, if *B* agrees to engage in communication with *A*, she must do so under the assumption that *A* doesn't take herself to be in the 'anything goes' situation.

Let us clarify and generalize the foregoing.

First, let us highlight the basic theoretical point brought out by the example. Suppose *B* takes a question *Q* to be a question under discussion. Then, we claim, this entails that *B* is accepting that *A* accepts that there is an 'objective' answer to *Q*—i.e., that *B* is accepting that *A* accepts that there is an answer to *Q* that she ought to accept. In other words, questions under discussions are (among other things) questions that we take to have an answer that we ought to endorse. We cannot take a question to be a question under discussion and at the same time think we may accept either answer to it.²⁵

Second, our example concerns a simple case where there is only one question under discussion. But the point generalizes. In every conversation, there is a set of questions under discussion $\{Q_1, Q_2, Q_3, \dots\}$.²⁶ The norms that are in the context set correspond to conjunctions of answers to all these questions. Now, suppose that, according to *A*, it is not the case that there is a norm *A* ought to accept. Then there is at least one question under discussion Q_i such that, according to *B*, it would make no sense for her to attempt to answer that question.

Third, as we hinted at already, the issue of whether it makes sense for a speaker *B* to engage in conversation with *A* should be relativized to what questions are under discussion. To see the point, it helps to consider the following objection.²⁷ Suppose we have three (mutually exclusive and jointly exhaustive) small norms: n_1 , n_2 , and n_3 . For vividness, let these norms concern three ethical stances about diet: n_1 allows omnivorism, n_2 mandates vegetarianism, n_3 mandates veganism. Now, suppose that Alice and Bob are discussing the ethics of dietary choices. Alice feels strongly about n_1 ; it's important to her that n_1 be ruled out. But she thinks that nothing settles whether one should adopt n_2 or n_3 . In her view, ethical rules 'run out' before settling which of n_2 or n_3 is right. In this case, one might worry, the context set of the conversation between

²⁵Strictly speaking, our example shows something slightly weaker. It shows that, *by attempting to settle one of the questions under discussion*, *B* is ruling out the possibility that *A* takes it to be permissible to adopt either answer. But, one might object, maybe there are questions that we take to be under discussion, but that we assume will never be settled. Won't our claim fail in this case? We think that this case can be safely ignored. In taking *Q* to be a question under discussion, one acts as if one believes one will get to address it. Much like in placing an item on the agenda one is acting as if one believes one will get to it during a meeting, even if one believes one may not get to it, in taking *Q* to be a question under discussion one is acting as if one of the participants in the conversation will attempt to answer *Q*.

²⁶For simplicity, we will assume that all questions under discussion are simple yes-no questions. Nothing in our argument hinges on this.

²⁷Thanks to an anonymous referee for pushing us to consider this case.

Alice and Bob fits our characterization of the ‘anything goes’ situation, since it is not the case that there is one particular norm that Alice ought not rule out. Yet Alice does not take herself to be allowed to accept any answer whatsoever to the questions under discussion. In this case, the worry goes, it might still make sense for Bob to engage in conversation with Alice, despite the fact that we are in an ‘anything goes’ scenario.

We agree with the objector that, in the case described, it makes sense for Bob to engage in conversation with Alice. But we deny that the questions under discussion are the ones they describe. If Alice really thinks that norms ‘run out’ before settling whether n_2 or n_3 is right, then the relevant question under discussion in the conversation is not which of n_1 , n_2 , or n_3 to adopt. Rather, the relevant question is which of n_1 , or (n_2 or n_3) is right. Recall from section 2.4: a question is under discussion just in case all speakers intend to engage in inquiry to settle it. But Alice does not intend to engage in inquiry to settle whether n_2 or n_3 is correct. She thinks ethical rules ‘run out’ before deciding which is correct. So the subquestion whether n_2 or n_3 is right is not on the agenda of the conversation.²⁸

Let us recap our basic point. A and B are in conversation. In the background, there are a number of questions they intend to discuss. B makes a claim that amounts to an answer to one of those questions—say ‘Omnivorism is ethically forbidden.’²⁹ For B ’s speech act to make sense, she must think that (a) it is under discussion whether omnivorism is wrong; and (b) A doesn’t think that she may accept either answer to the question whether omnivorism is forbidden. This will be our starting point in developing the rest of our argument.

Before moving on, let us make a final clarification. We are *not* claiming that, whenever an agent is in an ‘anything goes’ situation, it makes no sense to try to coordinate with her. Consider other instances of coordination, like choosing which side of the road to drive on. This case arguably does display the permissibility that we want to deny for the case of communication. Neither agents has a reason to pick one side rather than another and they take this to be so. Yet it still makes sense for them to make proposals about how to coordinate. So there is a point to coordinating activities even when agents take themselves not to be subject to any normative constraints.

What we *are* claiming is instead the following: in an ‘anything goes’ situation, coordination may not proceed *via assertions of declarative sentences*. An ‘anything goes’ situation allows for a gap between recognizing that a proposal to coordinate is legitimate and accepting that proposal. Our practice of making and accepting assertions doesn’t recognize this gap. To see the point, contrast the following two conversations:

²⁸Of course, Alice can, for the sake of the argument, act as if she thinks that ethical rules do settle whether n_2 or n_3 is correct. In this case, she *will* intend to engage in inquiry to settle whether n_2 or n_3 is correct, at least for the purposes of the conversation, and it makes sense to engage with Alice in conversation about the issue.

²⁹Or perhaps something weaker, i.e. a disjunction of answers to one or more questions.

A: I suggest that we drive on the right-hand side of the road.

B: Nah, no way.

C: The Brits drive on the left-hand side of the road.

D: Nah, no way.

Both *B* and *D* are rejecting what their interlocutor said. But notice a crucial difference. By rejecting *A*'s utterance, *B* is not taking issue with the appropriateness of *A*'s speech act. *B* can (and normally will) deem *A*'s speech act perfectly successful; *B* needs find no fault with *A*'s advancing the suggestion that they drive on the right. This is fully compatible with not accepting that suggestion. (Another way to see this: notice that *B* could express his rejection as: "Okay, you've made a perfectly fine suggestion. But I refuse to accept it.") But the situation is different for the case of *C* and *D*. *D*'s rejection amounts to a rejection of *D*'s speech act. *D* cannot think that there is nothing wrong with *C*'s speech act, and at the same time think it's okay not to update their belief as suggested by *C*. (Notice the awkwardness of saying: "Okay, you've made a perfectly fine assertion. But I refuse to accept it.") For the case of assertion, there is no gap between finding an assertion appropriate and taking up the content of that assertion.

Notice that we are *assuming* (on the basis of simple empirical observation) that assertions of moral claims work in this way.³⁰ This is one of the starting points of our argument. We are not trying to explain, on behalf of the expressivist, why this is so. Thus we're not trying to give a full account of communication on behalf of the expressivist. What we're doing is part of this enterprise, but we must leave the completion of the task to a different occasion.³¹

In summary: while coordination is possible (and indeed it does often happen) in an 'anything goes' situation, in such a situation it would make little sense to attempt coordination by way of utterances of declarative sentences.

³⁰Remember that we are using 'assertion' in the thin sense: our assumption is thus that this is the way utterances of declarative sentences work, at least in 'normal' cases.

³¹Let us just briefly flag a natural direction the expressivist might go to answer this challenge. She might argue that the speech act of assertion is connected to epistemic practices requiring us to provide shared reasons for backing assertions or for rejecting them. This idea connects with a pretty different line of thinking about assertion than the one we're using, namely the one developed in Brandom, 1994. We think that, in principle at least, these lines of thinking may be reconciled; they seem to investigate different and complementary aspects of assertion. Roughly: Stalnaker's account focuses on the informational aspect of assertion, and on how assertions affects the speakers' mental states in conversation; a Brandomian theory of assertion spells out the epistemic norms that govern the practice of assertion.

4.3 Establishing CAN GO WRONG

We have argued that, for it to make sense for B to engage in conversation with A , B must accept that A is not in an ‘anything goes’ situation:

$$\neg \text{Acc}_B(\text{Acc}_A \forall n(\text{May}(n \notin \text{Acc}_A))).$$

We now claim that this is enough to establish the following claim:

$$(A0) \quad \text{Acc}_B \text{Acc}_A (\exists n \text{Ought}(n \in \text{Acc}_A))$$

In plain English, this says that B accepts that A accepts that there is a norm A ought not rule out.³² Notice that (A0) says that B presupposes a *descriptive* proposition about A ’s attitudes: namely, the proposition that A accepts that there is one norm that she is not allowed to rule out.

Of course, it is not in general true that if B must not believe that p , B must therefore believe that $\neg p$. We argued that if B is convinced that, according to A , anything goes for her, then it makes no sense for her to engage in communication. But, wouldn’t it be enough for B to *leave open* the possibility that A takes herself to be subject to some constraints in what she accepts? Why think that (A0) is required for it to make sense for B to engage in conversation with A ?

It is here that the peculiar nature of the notion of acceptance becomes crucial to our argument. We grant that B may not *believe* that it’s not true that A accepts that anything goes for A . Nevertheless, for the purposes of conversation, she must take that for granted. To make an assertion is not to *try* to make a demand that the hearer changes her attitude. Rather, it just is to make a demand to that effect. By making an assertion, the speaker represents herself as accepting that it’s not true that anything goes. Hence she takes that for granted for the purposes of the conversation.

Three intermediate premises. Once we have (A0), it’s easy to get a further claim that will be useful. If A is rational, she will see the reasoning explained above and realize that (A0) is required for communication to function.³³ Hence she will accept (A0), which gives us the following:

$$(A1) \quad \text{Acc}_A \text{Acc}_B \text{Acc}_A (\exists n \text{Ought}(n \in \text{Acc}_A)).$$

³²There are modes-of-presentation-kind issues with this formulation. What we want to say is that B accepts that A accepts that there is a norm that *she herself*, individuated in the self-locating way (see Lewis, 1979a; Perry, 1979, among many), ought not rule out. These issues are orthogonal to our main concern here and can be put to the side.

³³Of course, this is not to say that A will have explicit knowledge of (A0), or that she would be able to report that she accepts (A0) for the purposes of the conversation. Knowledge and acceptance of the principles discussed in this section may be implicit.

Given these two claims, we can establish that each speaker accepts that there is a norm that she ought not rule out. In symbols, this involves eliminating two iterations of the ‘Acc’ operator from (A1):

$$(A2) \quad \text{Acc}_A(\exists n \text{ Ought}(n \in \text{Acc}_A)).$$

To see that, suppose that *B* accepts that *A* accepts φ , and that *A* accepts that *B* accepts that *A* accepts that φ , for some proposition (descriptive or normative) φ :

$$(i) \quad \text{Acc}_B \text{Acc}_A \varphi$$

$$(ii) \quad \text{Acc}_A \text{Acc}_B \text{Acc}_A \varphi$$

Now, assume that, despite this, *A* doesn’t accept φ , i.e. $\neg \text{Acc}_A \varphi$. Then *A* has a choice. She can either point out that she does not accept φ , or she can act as if she accepts it for the purposes of the conversation. If she goes for the first option, *B* will stop accepting that *A* accepts φ . If she goes for the second option, she will actually come to accept φ for the purposes of the conversation. (Notice that here, once more, the particular nature of acceptance and its divergence from belief is crucial to make the point.)

Now, our claims (A0) and (A1) are just instances of schemas (i) and (ii). We established above that (A0) and (A1) hold in any ‘normal’ conversation. Hence, in any normal conversation, speakers are faced, at any point, with the alternative of denying that they accept that there is a norm they ought not to rule out, or going along with it. If they make the former choice, then they will essentially opt out of the conversation. In that case, in fact, (A0) will no longer hold and communication will lose its point. Hence, if the conversation goes on, then they effectively accept that there is a norm they ought not rule out.

In symbols, by running the reasoning for both speakers, we get the following two claims:

$$(A2) \quad \text{Acc}_A(\exists n \text{ Ought}(n \in \text{Acc}_A)).$$

$$(A3) \quad \text{Acc}_B(\exists n \text{ Ought}(n \in \text{Acc}_B)).$$

These say that each participant in a conversation accepts that there is a norm that she ought not rule out. To get CAN GO WRONG from here, we still have to show that they all accept that there is a *unique* norm that they *both* ought not rule out.

The final step. The argument for this conclusion is not dissimilar from the argument we gave for (A0). Suppose that it’s not true that speakers accept that there is a norm they both ought to accept. For example, suppose that *B* doesn’t accept that. By the previous argument, we know that *B* accepts that there is a norm that she ought not rule out: call that norm ‘ n^* ’. According to *B*, it might be that *A* is allowed to rule out n^* . This

allows for a situation of the following kind. Suppose that *A* utters a normative claim: for example, ‘Tax evasion is wrong’. Now, *B* may accept that *A*’s assertion is correct in a number of respects: it is sincere, it doesn’t violate any constraints on *A*’s attitudes, etc. Nevertheless, since *B* doesn’t accept that she and her interlocutor are coordinating on the same norm, she has no reason to assent to *A*’s assertion. It might be that the proposition expressed by ‘Tax evasion is wrong’ is okay to accept for *A*, but not for *B*—i.e., it might rule out n^* , the norm that *B* ought not rule out.

In this situation, it is not rational for *B* to engage in communication with *A*. If she did, she would risk being in violation of the normative requirement that she takes herself to be subject to (via (A3)). Hence, in order for *B* to engage in communication with *A*, she will accept that there is a unique norm that both *A* and *B* ought not rule out. Of course, a parallel conclusion holds for *A*.

Taking stock, this means that all speakers involved in conversation will accept the following:

$$(A4) \quad (\exists n)(\forall s)\text{Ought}(n \in Acc_s)$$

From here, it’s a short step to claim that (A4) has common ground status. After all, we did not need any specific assumptions about what speakers happened to accept in order to establish (A4). Thus, if common belief in minimal rationality is in place, it will be common belief, at least tacitly, that (A4) is accepted by all participants in the conversation. This gives us exactly CAN GO WRONG, which we repeat below:

CAN GO WRONG. The following is common ground in any conversation:

$$(\exists n)(\forall s)\text{Ought}(n \in Acc_s)$$

Summary. It’s time to take stock. We have bootstrapped our way into our conclusion from a number of intermediate premises. We have claimed, first, that while engaging in conversation each speaker must assume that ‘not anything goes’ for her interlocutor. From here, we have argued, exploiting the peculiar public nature of acceptance, that each speaker must assume that ‘not anything goes’ for her own attitudes as well. Finally, we have argued that, if speakers have to trust each other, they must assume that ‘not anything goes’ for them in the same way—essentially, that the normative constraints applying to all of their attitudes are the same.

4.4 Establishing UNIQUENESS

In order to establish CONVERGENCE, we need one more claim. This is the claim that speakers presuppose that their acceptance states ought to be shrunk to a unique point:

UNIQUENESS. The following is common ground in any conversation:

$$\text{Ought}((\exists n)(\forall s)Acc_s = \{n\}).$$

UNIQUENESS may seem more controversial than it is. At first sight, you might worry that it enforces a requirement of objectivity that is even stronger than that posed by CAN GO WRONG. But notice the difference in the scope of the operators. UNIQUENESS doesn't say that there is a norm on which acceptances ought to converge; it just says that speakers ought to converge on some norm or other. Hence UNIQUENESS amounts to no more than the claim that, when engaged in conversation, we take for granted that we ought to agree on one among a number of live alternatives (and that this is common belief between us). But now, at least given our working picture of conversation, it is a truism that the point of conversation is to achieve consensus as to which of the live alternatives is the one we ought to coordinate on. Hence we take UNIQUENESS to be uncontroversial.

Once again, it is worth recalling that we are working with small worlds and small norms. It would be unrealistic to assume that ordinary speakers take up the task of narrowing open possibilities to a point in logical space. But we are not making this assumption. On the small worlds construal, UNIQUENESS just amounts to the requirement that all questions under discussion receive an answer. This much seems perfectly acceptable. By raising a question, a speaker seems to demand that participants in the conversation make a choice between the available answers.

To be sure, speakers normally don't hold that all uncertainty ought to be eliminated at any cost. They can opt out of conversation for a number of reasons—they're uninterested, or it's too demanding, or it's clear that agreement won't be reached. But this is compatible with UNIQUENESS. We are claiming that UNIQUENESS is something that speakers hold *as long as they are engaged in conversation*. In other words, as long as they think that there is a point to engaging in conversation, they must think that they ought to converge on some live possibility. Converging on some live possibility is just what the point of conversation is. Speakers may give up on UNIQUENESS, but by doing so they give up on the very point of staying in a conversation.

4.5 Convergence

We have argued for the following two claims:

CAN GO WRONG. The following is common ground in any conversation:

$$(\exists n)(\forall s)\text{Ought}(n \in Acc_s)$$

UNIQUENESS. The following is common ground in any conversation:

$$\text{Ought}((\exists n)(\forall s)Acc_s = \{n\}).$$

These two claims entail our wanted conclusion:

CONVERGENCE. In any normal conversation, the following is common ground:

$$(\exists n)(\forall s)\text{Ought}(\{n\} = Acc_s).$$

The step from CAN GO WRONG and UNIQUENESS to CONVERGENCE is straightforward. CAN GO WRONG establishes that there is one norm that ought not be ruled out from the context set. UNIQUENESS establishes that the context set ought to be reduced to some norm or other. The only way to satisfy the two desiderata is that the context set be shrunk exactly to the one norm that ought not be ruled out.³⁴

4.6 Convergence and Error Theories

We have argued that the following claim is presupposed in any normal conversation:

$$(\exists n)(\forall s)\text{Ought}(\{n\} = Acc_s).$$

Of course, this is not to say that this claim is true. Indeed, it's not clear how the question of truth and falsity of normative claims should be handled by the expressivist in the metalanguage.

One consequence of this is worth highlighting. The model of communication that results from our argument doesn't require any kind of pretense or error theory. Since the crucial presupposition is normative, speakers are not required to make an assumption that may be false by the expressivist's own lights (something in the ballpark of there being normative facts) to proceed in conversation. The assumption they must make is functionally analogous. But, crucially, it doesn't commit them to some false claim about what the world is like.

We take this to be a significant advantage of our proposal. To see why, let's compare it briefly with a somewhat similar picture of communication for nonfactual discourse proposed by Andy Egan (2007; 2010), in the context of defending a form of relativism. Egan endorses a form of truth relativism on which the contents of certain assertions—for example, claims about personal taste—are centered world propositions, i.e. sets of triples of a world, a time, and an individual. On this picture, the content of

³⁴Here is a more detailed proof. Assume, with CAN GO WRONG, $(\exists n)(\forall s)\text{Ought}(n \in Acc_s)$. Let n^* be such an n ; we get that $(\forall s)\text{Ought}(n^* \in Acc_s)$. Then assume, with UNIQUENESS, that $\text{Ought}((\exists n)(\forall s)Acc_s = \{n\})$. Finally, for *reductio* assume that CONVERGENCE fails, i.e. assume that $(\forall n)(\exists s)\text{May}\neg(\{n\} = Acc_s)$. Instantiating the last claim with n^* , we get that $(\exists s)\text{May}\neg(\{n^*\} = Acc_s)$. Let s^* be the relevant speaker: we get $\text{May}\neg(\{n^*\} = Acc_{s^*})$. Given CAN GO WRONG, this means that s^* 's acceptance state can permissibly be a (proper) superset of $\{n^*\}$. But UNIQUENESS requires that $\text{Ought}((\exists n)Acc_{s^*} = \{n\})$, i.e. that s^* 's acceptance state be a singleton. Contradiction.

(5) Eggplants are tasty.

is the set of triples $\langle w, t, i \rangle$ such that i has (at t in w) the disposition to like eggplants (Egan, 2010, p. 259ff).

Since contents are more fine-grained than sets of worlds, Egan has a problem analogous to that of the expressivist. He must explain how assertions of claims about taste work, given that he can't appeal to the standard idea of locating an actual point. He does this by imposing a constraint on the felicity of assertions: asserting a centered-world proposition requires presupposing that all speakers in the conversation are similar in relevant respects. For example, a felicitous assertion of (5) requires that all speakers presuppose that they're disposed to have a similar response of enjoyment or distaste towards eggplants.

This proposal seems to run into trouble rather quickly. After all, (5) seems assertable even in cases where it's entirely manifest that one's audience doesn't share one's dispositions ("What? You don't like eggplants? You're crazy! Eggplants are just *so* tasty!").³⁵ An account that requires a factual presupposition of similarity to be in place is hard-pressed to explain cases like this, since this presupposition flatly contradicts information that is built into the context set.³⁶

To be sure, this does not show that there is no account in terms of a presupposition of similarity that could be made to work.³⁷ But even if such an account could yield the right predictions, we are still being forced to make controversial assumptions about the level of conceptual sophistication one would need in order to believe that eggplants are tasty. In contrast, our model has no problems with cases like this. On our model, just engaging in conversation is enough to bring about the presupposition that there ought to be a point of convergence between speakers. While it might be evident that convergence is hard to achieve in some cases, this doesn't contradict other bits of common ground information.

³⁵Note that this worry is different from that of whether disputes over (5) are ever worthwhile. Egan rightly points out that one can be mistaken about whether one is disposed to enjoy eggplants. Suppose that Alice utters (5), and Barbara utters its negation. On Egan's view, Alice can take Barbara to have been sincere while still thinking she is wrong—after all, it might be that Barbara is mistaken about her own disposition to enjoy eggplants. Our point is that, on Egan's picture, Alice cannot explicitly recognize that Barbara lacks that disposition and continue to insist that eggplants really *are* tasty. Thanks to an anonymous referee for pushing us to clarify this point.

³⁶As Lewis, 1979a points out, it is customary to bring about a presupposition in the context set by making an assertion that requires it. This is the so-called phenomenon of *presupposition accommodation*. Note however that, in this case, the presupposition could not just be added to the context set, for it contradicts other information that is already in the context set. While there may be cases where a speaker triggers accommodation of information that conflicts with the context set, it seems odd to utter in one breath two sentences one of which presupposes something that is incompatible with the other. On Egan's picture, an utterance of 'I see you do not like eggplants, but they really are tasty' would be like an utterance of 'There is no King of France, but the King of France is bald'. Thanks here to an anonymous referee.

³⁷For a different proposal, Cf. Egan, 2010, p. 270f.

5 Conclusion

Expressivism about normative discourse runs into a problem about communication. The expressivist wants to see communication as a form of coordination in attitudes among speakers. But it is not clear how the expressivist can explain why the communication of normative claims takes the form that it does. The descriptivist model of communication makes an obvious appeal to truth, and it's unclear how the expressivist can do without it.

We argued that, just by assuming certain basic facts about communication and the rationality of the speakers, we can answer this challenge. The expressivist can think of normative conversation as a joint attempt at ruling out possibilities—much like the descriptivist does—where these are systems of norms that are complete with regards to all questions under discussion. The similarity between normative and non-normative conversation is due to the fact that, when engaging in normative conversation—at least in so far as these are adequately modeled in a broadly Stalnakerian picture—speakers need to assume that there is a unique system of norm on which their attitudes ought to converge. This assumption works as the normative counterpart of the presupposition that there is an actual world that speakers are trying to locate in conversation.

Let us close with some more speculative remarks. The problem of explaining how communication works is part of a more general problem for the expressivist, i.e., the problem of accounting for the apparent objectivity of moral thought and talk. It's a vexed question exactly what 'objectivity' means in this context (see e.g. Rosen, 1994). Following Gibbard, let's just say that there is a sense in which, by making a normative statement, a speaker seems to be claiming an authority of some sort. She's not merely making a suggestion as to how conversational participants could achieve coordination. Rather, she's claiming that her opinion is right in a stronger, intersubjective sense.³⁸ One major issue for the expressivist is doing justice to this idea.

Our conclusions in this paper are certainly not enough to address this issue. But we think we have taken some steps in the right direction. If our argument is sound, we've shown that engaging with anyone's normative claims requires assuming the existence of a kind of intersubjective normative standard. Part of what it is to engage in normative conversation with an agent is to take her claims to apply some pressure on one's own normative views. This pressure comes not from a factual assumption of a common normative outlook, but rather from the normative structure of conversation itself.

What we have not shown, of course, is how this kind of intersubjectivity can stretch beyond the boundaries of conversations. For example, our account is silent on why Anne and Zoe, who have never met face to face and never been in a conversation together, should take themselves to be subjected to the same normative standard. For all we know, nothing that we've said helps answer this question. But our general strategy—

³⁸Cf. Gibbard, 1990, p. 155.

i.e., deriving a kind of objectivity from the structure of our social practices—might be worth exploring as a way of producing a more general vindication of objectivity on behalf of the expressivist.³⁹

References

- Brandon, R. (1994). *Making it Explicit*. Cambridge, MA: Harvard University Press.
- Dorr, C. (2002). “Non-cognitivism and Wishful Thinking”. *Noûs* 36.1, pp. 97–103.
- Dreier, J. (1999). “Transforming Expressivism”. *Noûs* 33.4, pp. 558–572.
- Dummett, M. (1959). “Truth”. *Proceedings of the Aristotelian Society* 59, pp. 141–162.
- (1981). *Frege: Philosophy of Language*. Second. Cambridge, Mass.: Harvard University Press.
- Egan, A. (2007). “Epistemic Modals, Relativism and Assertion”. *Philosophical Studies* 133.1, pp. 1–22.
- (2010). “Disputing about Taste”. In: *Disagreement*. Ed. by T. A. Warfield and R. Feldman. Oxford: Oxford University Press, pp. 247–286.
- Geach, P. (1960). “Ascriptivism”. *Philosophical Review* 69.2, pp. 221–225.
- (1965). “Assertion”. *Philosophical Review* 74.4, pp. 449–465.
- Gibbard, A. (1990). *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Cambridge, Mass.: Harvard University Press.
- Groenendijk, J. and M. Stokhof (1984). “Studies in the Semantics of Questions and the Pragmatics of Answers”. PhD thesis. Amsterdam: University of Amsterdam.
- Hamblin, C. L. (1958). “Questions”. *Australasian Journal of Philosophy* 36.3, pp. 159–168.
- (1973). “Questions in Montague English”. *Foundations of Language* 10.1, pp. 41–53.
- Hulstijn, J. (1997). “Structured information states: Raising and resolving issues”. In: *Proceedings of MunDial*. Vol. 97, pp. 99–117.
- Joyce, J. M. (1999). *The Foundations of Causal Decision Theory*. New York: Cambridge Univ Press.
- Lewis, D. (1979a). “Attitudes *De Dicto* and *De Se*”. *Philosophical Review* 88.4, pp. 513–543.
- (1979b). “Scorekeeping in a language game”. *Journal of Philosophical Logic* 8.1, pp. 339–359.
- (1980). “Index, Context, and Content”. In: *Philosophy and Grammar*. Ed. by S. Kanger and S. Öhman. Dordrecht: Reidel, pp. 79–100. Reprinted in Lewis 1998, pp. 21–44.
- (1988). “Statements partly about observation”. *Philosophical Papers* 17. Reprinted in Lewis, 1998., pp. 1–31.
- (1998). *Papers in Philosophical Logic*. Cambridge: Cambridge University Press.
- Ninan, D. (2010). “Semantics and the objects of assertion”. *Linguistics and Philosophy* 33.5, pp. 355–380.
- Pérez Carballo, A. (2012). “Negation, Expressivism, and Intentionality”. Unpublished typescript, University of Massachusetts, Amherst.

³⁹Thanks to audiences at Harvard, MIT, the Australian National University, and the University of Sydney, to two anonymous referees for *Ethics*, as well as to the journal’s associate editors for helpful comments and discussion. Special thanks to Fabrizio Cariani, John MacFarlane, David Plukett, Agustín Rayo, Bob Stalnaker, and Judy Thomson for feedback at many different stages of the project.

- Perry, J. (1979). "The Problem of the Essential Indexical". *Noûs* 13.1, pp. 3–21.
- Price, H. (2011). "Expressivism for Two Voices". In: *Pragmatism, Science, and Naturalism*. Ed. by J. Knowles and H. Rydenfelt. New York: Peter Lang, pp. 87–114.
- Rabern, B. (2012). "Against the identification of assertoric content with compositional value". *Synthese* 189.1, pp. 75–96.
- Roberts, C. (2012). "Information structure in discourse: Towards an integrated formal theory of pragmatics". *Semantics and Pragmatics* 5.6, pp. 1–69.
- Rosen, G. (1994). "Objectivity and Modern Idealism: What is the Question?" In: *Philosophy in Mind*. Ed. by M. Michael and J. O'Leary-Hawthorne. Dordrecht: Kluwer.
- Rothschild, D. (2012). "Expressing Credences". *Proceedings of the Aristotelian Society* 112.1pt1, pp. 99–114.
- Savage, L. J. (1972). *The Foundations of Statistics*. Second. New York: Dover.
- Schroeder, M. (2008a). *Being For: Evaluating the Semantic Program of Expressivism*. Oxford: Oxford University Press.
- (2008b). "How expressivists can and should solve their problem with negation". *Noûs* 42.4, pp. 573–599.
- (2010). *Noncognitivism in Ethics*. Oxford: Routledge.
- Stalnaker, R. C. (1973). "Presuppositions". *Journal of Philosophical Logic* 2.4, pp. 447–457. Reprinted in Stalnaker, 1999.
- (1978). "Assertion". In: *Syntax and Semantics*. Ed. by P. Cole. Vol. 9. New York: Academic Press, pp. 315–322. Reprinted in Stalnaker, 1999.
- (1981). "Indexical Belief". *Synthese* 49.1. Reprinted in: Stalnaker, 1999, pp. 130–149. Page numbers refer to the reprinted version., pp. 129–151.
- (1999). *Context and Content*. Oxford: Oxford University Press.
- (2002). "Common Ground". *Linguistics and Philosophy* 25.5, pp. 701–721.
- (2014). *Context*. Oxford: Oxford University Press.
- Swanson, E. (2011). "How Not to Theorize about the Language of Subjective Uncertainty". In: *Epistemic Modality*. Ed. by A. Egan and B. Weatherson. Oxford: Oxford University Press, pp. 249–269.
- Veltman, F. (1996). "Defaults in update semantics". *Journal of Philosophical Logic* 25.3, pp. 221–261.
- Williamson, T. (1996). "Knowing and Asserting". *Philosophical Review*, pp. 489–523.
- (2000). *Knowledge and its Limits*. Oxford: Oxford University Press.
- Yalcin, S. (2011). "Nonfactualism about Epistemic Modality". In: *Epistemic Modality*. Ed. by A. Egan and B. Weatherson. Oxford University Press, pp. 295–334.
- (2012). "Bayesian Expressivism". *Proceedings of the Aristotelian Society* 112.2, pp. 123–160.