

Conventions Without Knowledge of Conformity

Megan Henricks Stotts

This version of the article has been accepted for publication in *Philosophical Studies* after peer review but is not the Version of Record and does not reflect post-acceptance improvements, or any corrections. The Version of Record is available online at: <https://doi.org/10.1007/s11098-023-01957-z>. Use of this Accepted Version is subject to the publisher's Accepted Manuscript terms of use <https://www.springernature.com/gp/open-research/policies/accepted-manuscript-terms>.

1. Introduction

Many aspects of David Lewis's (1969/2002, 1975) account of conventions have received careful and compelling criticism in the years since the account's publication.¹ But one aspect of Lewis's account has not attracted controversy, and his critics have often retained it in various forms in their own accounts of conventions: Lewis's knowledge of conformity requirement. As I will use the term, a *knowledge of conformity requirement* mandates that members of a group in which doing *A* is conventional must know that they and/or others do *A*. Importantly, knowledge of conformity in this sense does not require participants to know that doing *A* is conventional, nor does it require them to know that there are alternatives to doing *A*.² Although knowledge of conformity requirements may seem innocuous, I will argue that they are misguided, and then I'll build an account of conventions that does not require knowledge of conformity.

After establishing (in Section 2) that Lewis and others require knowledge of conformity, I'll argue (in Section 3) that such requirements, even when softened to require only potential or tacit knowledge, wrongly exclude a key class of paradigmatic conventions that permeate our lives. Most saliently, this class includes some central semantic conventions. In Section 4, I'll show that Ruth

¹ See, e.g., Burge, 1975; Gilbert, 1989; Miller, 2001; Davis, 2003; and Marmor, 2009.

² Lewis (1969/2002, 1975) does require both of these other kinds of knowledge, in addition to knowledge of conformity (pp. 61, pp. 5–6). But these other requirements, unlike his more minimal knowledge of conformity requirement, have already received considerable criticism, such as from some of the authors cited in note 1.

Garrett Millikan's (2005, 2014) account, which does not require knowledge of conformity, can accommodate these conventions. However, as I'll argue, Millikan's account is marred by her mistaken claim that most linguistic conventions involve coordination, and by complications related to a key element of her account: reproduction. So, in Section 5, I'll offer an amended notion of copying to replace Millikan's notion of reproduction and then propose an account of conventions that is informed by Millikan's insights but which avoids her account's problems. More specifically, I'll contend that doing *A* is conventional in groups in which there is widespread, interconnected copying of doing *A* as a way of doing something further, *B*, where there are other equally good and accessible ways to do *B*.

2. Knowledge of conformity requirements

The first step is to establish that Lewis, and certain authors who have followed him, really does require knowledge of conformity within the population or group in which a convention obtains.

Lewis (1975) offers the following account of conventions:

a regularity *R*, in action or in action and belief, is a *convention* in a population *P* if and only if, within *P*, the following six conditions hold. (Or at least they almost hold. A few exceptions to the "everyone"s can be tolerated.)

- (1) Everyone conforms to *R*.
- (2) Everyone believes that the others conform to *R*.
- (3) This belief that the others conform to *R* gives everyone a good and decisive reason to conform to *R* himself...
- (4) There is a general preference for general conformity to *R* rather than slightly-less-than-general conformity—in particular, rather than conformity by all but any one...
- (5) ...There is at least one alternative *R'* such that [*R'* would satisfy (3) and (4), and]... *R'* could have perpetuated itself as a convention instead of *R*...
- (6) Finally, the various facts listed in conditions (1) to (5) are matters of *common* (or *mutual*) knowledge (pp. 5–6).

Lewis's knowledge of conformity requirement comes from the application of condition (6) to condition (1). It must be common knowledge that everyone conforms to the regularity.

For now, we'll take Lewis's knowledge of conformity requirement at face value, but I do want to note that Lewis (1969/2002, 1975) proposes two different ways of softening his claims about

common knowledge: allowing the common knowledge to be “merely potential,” and allowing it to be “tacit” (pp. 63–64; p. 6). We’ll discuss these strategies in Section 3.

Before showing how knowledge of conformity requirements have manifested in other authors, it’s important to consider what may have motivated Lewis to require knowledge of conformity in the first place. Lewis (1969/2002) developed his account of conventions to vindicate the claim that language is conventional. To succeed, he needed an account of conventions that did not presuppose that participants can communicate with each other in advance of setting up conventions (p. 2). A population that does not yet have a language confronts an embarrassment of riches: a vast number of possible languages would work equally well, if only they could somehow settle on one of them (pp. 7–8). What the population needs, according to Lewis, is to get to the point where each member has some basis for knowing which of the many possible languages other members will use in the future. For Lewis, one key part of getting to that point is knowing that others in the group do, in fact, conform to some particular regularity in their communicative practices (pp. 36–37).³ Knowledge that one conforms oneself is baked in as well, perhaps because it just seemed obvious that once one was aware of others’ conformity, one would be aware of one’s own as well.

Seumas Miller (2001) includes a similar requirement. The following is his account of conventions:

There is a convention to x in the recurring situation s among agents A, B, C , and so on, if and only if:

- 1) each has the procedure to x in s ;
- 2) each has the procedure to x in s if and only if the others have the procedure to x in s ;
- 3) each has the collective end e , and each x ing in every instance of s realises e ;
- 4) each mutually truly believes that everyone has the procedure to x in s (p. 118).

For Miller, having a procedure to x in s involves “the automatic repetition of an action in a recurring situation,” which implies actual conformity (p. 25). And condition (4) requires mutual true belief that

³ Margaret Gilbert (1989) provides a compelling criticism of this line of reasoning (pp. 330–338). But at this stage we are concerned just with what motivated Lewis, so I won’t criticize his view on these grounds.

everyone has such a procedure, and thereby, mutual true belief that everyone conforms. Miller defines mutual true belief as follows: “*A* and *B* mutually truly believe that *p*, if and only if: (1) *p*; (2) *A* believes that *p*, and *B* believes that *p*; (3) *A* believes that *B* believes that *p*, and *B* believes that *A* believes that *p*, and so on” (p. 59). Lewis (1975) endorses a similar definition of common knowledge (p. 6). Because the distinction between knowledge and true belief is unimportant for our purposes, I will just use the term ‘common knowledge’ going forward, to cover Miller’s and Lewis’s notions. So, we’ve seen that Miller, too, requires common knowledge of conformity within the relevant group.

Andrei Marmor (2009) also requires a kind of knowledge of conformity. His account of conventions is as follows:

A rule, *R*, is conventional, if and only if all the following conditions obtain:

1. There is a group of people, a population, *P*, that normally follow *R* in circumstances *C*.
2. There is a reason, or a combination of reasons, call it *A*, for members of *P* to follow *R* in circumstances *C*.
3. There is at least one other potential rule, *S*, that if members of *P* had actually followed in circumstances *C*, then *A* would have been a sufficient reason for members of *P* to follow *S* instead of *R* in circumstances *C*.... (p. 2).

The knowledge of conformity requirement is considerably subtler here, but it is still present, built into Marmor’s understanding of rules. Marmor writes: “The idea of following a rule ... normally entails that the agent regards the rule as binding under the circumstances, which would normally entail that the agent must be aware of the fact that he or she is following a rule” (pp. 6–7). So, by treating conventions as a species of rules, Marmor implicitly requires participants to be aware of following the relevant rule. Each group member need only have knowledge of their *own* conformity, and this knowledge is just individual knowledge rather than common knowledge, but we will see that the same problems arise.

We’ll also discuss Margaret Gilbert’s (1989) account of conventions, which departs from Lewis’s account in deep ways but nonetheless includes a closely related knowledge requirement. According to Gilbert, a social convention is “a jointly accepted principle of action, a group fiat with

respect to how one is to act in certain situations” (p. 377). She requires that individual members of a population in which a convention obtains “know that they jointly accept” the relevant principle (p. 373). Because semantic conventions will be our main focus below, I’ll note that for Gilbert, semantic conventions are a bit different from what she calls ‘social conventions.’ But she makes clear that semantic conventions, too, come from a group fiat (p. 386), and thus they would presumably share the feature of knowledge of joint acceptance. So, although Gilbert does not require knowledge of conformity (or even conformity (p. 399)), she does require knowledge of joint acceptance of a principle.⁴

Table 1 summarizes the three main kinds of knowledge of conformity requirements that we will target in the next section.

Table 1: Kinds of knowledge of conformity requirements

		Whose conformity?	
		Everyone in the group/population	Each individual’s own
Kind of knowledge	Common knowledge	Lewis and Miller	
	Ordinary individual knowledge	None of our authors espouse this view, but we will entertain it.	Marmor

In addition to the various possible knowledge of conformity requirements in the table, we’ll discuss Gilbert’s similar requirement of knowledge of joint acceptance.

⁴ In fact, it seems that Gilbert (1989) requires not just individual, but common knowledge (pp. 203, 373). I will leave this detail aside, as my concerns about Gilbert’s knowledge requirement apply, *a fortiori*, to a version requiring common knowledge.

3. Against requiring knowledge of conformity

In this section, I'll advance a criticism that applies to every knowledge of conformity requirement in Table 1. Then, I'll extend the criticism to apply to Gilbert.

Demonstrating the full force of the criticism will be a complex task, but the criticism itself is simple: the problem with knowledge of conformity requirements is that there are many paradigmatic conventions where such knowledge is clearly absent. Consider, first, semantic conventions. On one hand, there are some semantic conventions with respect to which people probably *do* have knowledge, and even common knowledge, that they and others conform. For instance, it is at least somewhat plausible that English-speakers know what practice they all follow with 'I,' and even know that others know this, *etc.* But on the other hand, not all semantic conventions are so transparent. There are many semantic conventions that cannot be described without special training, making it implausible that people's conformity to them is common knowledge.

Consider, for instance, the word 'of.' If asked what their practice is with that word, a relatively introspective English-speaker might say that they use it to express ownership. Then they might pause before adding, "or some other close association between two things." I doubt many speakers would be much more specific than that. Similar considerations apply to words such as 'might,' 'but,' 'own,' 'if,' 'that,' 'ought,' and the German word 'doch.' And even certain nouns, such as 'element,' 'bunch,' or 'account,' or relatively simple verbs such as 'strain' or 'charge,' seem likely to pose similar challenges. That many native speakers cannot accurately describe their own and others' practice with these words suggests that they do not know that they or others conform to that practice, whether in the sense of individual knowledge or common knowledge.⁵

⁵ Wayne Davis (2003) objects to Lewis in a similar manner, providing some other kinds of examples of linguistic conventions without knowledge of conformity (pp. 210, 227).

In addition to semantic conventions, there are other subtle conventions to which we conform without realizing that we or others do so. For instance, conventions likely shape how we hold our bodies when waiting in line, sitting at a table in a restaurant, or browsing items for sale at a store; how we hold our personal belongings when out and about in public; how much we smile at others; and the extent to which we touch our own hair and skin when in the company of others with whom we have varying levels of familiarity. I am unable to describe my own practices with respect to any of these activities, and I doubt I am unique in this deficiency.⁶

These non-linguistic examples show that the problem is not restricted to semantic conventions, but really, for most of our authors, semantic conventions make the point best. Lewis, Miller, or Marmor could object that they didn't intend their views to cover the subtle social practices mentioned in the previous paragraph, but they all make it clear that they take their accounts to apply to semantic conventions (Lewis, 1975; Miller, 2001, p. 119; Marmor, 2009, p. 82).⁷

At this stage, one might object that Lewis would reject my examples of semantic conventions. On Lewis's (1975) view, semantic conventions are conventions of truthfulness and trust in whole languages, where languages are abstract functions from sentences to meanings. The examples I've offered operate at the level of individual words, and my criticism may seem unfair.

But as I see it, shifting to word-level semantic conventions actually benefits Lewis. The concerns I've raised about people's lack of common knowledge of conformity to the conventions for using 'of' or 'might' are small in comparison to the concerns I would want to raise about the idea that it is common knowledge within various populations that members follow regularities of truthfulness and trust with respect to abstract functions from sentences to meanings.

⁶ Millikan (2005) offers the example of standing distances (*i.e.*, the amount of space people maintain between themselves and nearby others) (p. 5). I do not list that example because empirical investigation has left the question of whether standing distances are conventional quite uncertain (Hayduk, 1983; Sorokowska *et al.*, 2017).

⁷ This claim may seem odd with respect to Marmor (2009), given that he says that semantics is by and large not conventional (p. 83). But, he does make it clear (in the passage cited in the main text) that he takes the sound-sense connection to be conventional, and that is our focus here.

Nonetheless, one might wonder whether my concerns could be avoided by moving to a simpler sentence-focused approach than Lewis's, where conventions just straightforwardly assign meanings to sentences. Problems with this move arise quickly, however, with respect to sentences that are never uttered. There will not be a regularity in how people utter them (as Lewis or Miller would require), nor do people have a rule that they normally follow in certain circumstances (*à la* Marmor). Even if we shifted to Gilbert's approach, which does not require any actual conformity, the idea that people have jointly accepted a group fiat about each sentence, including incredibly long ones that will never be uttered, just isn't plausible. It seems, then, that in order to satisfy all aspects of the accounts of conventions we've discussed, a sentence-centered semantics would have to include something akin to Lewis's move of positing a single, abstract function assigning meanings to sentences, and again, that kind of approach raises even more daunting concerns with respect to knowledge of conformity. So, I don't think it's problematic to focus on word-level semantic conventions.

As mentioned in Section 2, Lewis offers two different strategies that could apply to the cases I've listed as challenges to knowledge of conformity requirements. First, he allows that the common knowledge in question may be "merely potential: knowledge that would be available if one bothered to think hard enough" (Lewis, 1975, p. 6). For Lewis, although population members do not actually have common knowledge that everyone conforms to a particular practice involving 'might,' they do have that knowledge potentially, in the sense that they could gain it if they "bothered to think hard enough." Marmor (2009), too, ultimately requires just that "there is always the *potential* of awareness that in complying with a convention one follows a rule" (p. 7).

This kind of response might work with respect to some of our non-linguistic examples. Perhaps a bit of hard thinking *would* give us knowledge of our practices when we sit at a restaurant table. But, there is no amount of just hard thinking that would allow the average language-user to correctly describe English-speakers' practice with 'might.' Rather, people need specialized training,

and *then* hard thinking. If the fact that I could become aware of something if I received specialized training and then thought very hard meant that I had potential knowledge of that thing, then I potentially know nearly everything other humans know, plus other things we could eventually discover.

Of course, this is true, in a loose sense. On that interpretation, knowledge of conformity requirements would amount to something like the following: the fact that I and/or others conform, like most other facts, is knowable to the human mind, after adequate training. But this obvious truth is not something that needs to be added to an account of conventions. So, when a knowledge of conformity requirement is included, its presence implies something beyond that obvious truth: specifically, it implies that hard thinking alone could get people to knowledge (or even common knowledge) that they and/or others conform. And as we've discussed, this just isn't the case for certain semantic conventions.

Stephen Schiffer's (1972) account of conventions includes a requirement closely related to knowledge of conformity, and he offers a related strategy we should briefly consider (p. 154). For semantic conventions specifically, instead of following Lewis in requiring potential common knowledge of actual conformity, Schiffer requires *actual* common knowledge of *potential* conformity. That is, he requires that it be common knowledge within the group "that if someone utters x (in such-and-such circumstances), then he will mean such-and-such" (p. 155). This does not require that anyone has actually conformed; it requires only that it be common knowledge that if someone *were* to utter a given expression, they would mean such-and-such. But this strategy will not help. The "such-and-such" is not something most language-users are able to formulate, for a wide range of semantic conventions, and so they are no more able to grasp it when the conformity is potential than when it is actual.

Lewis's (1969/2002) other strategy is to say that people's knowledge can be "tacit" and "irremediably nonverbal" (pp. 63–64). Again, Marmor (2009) makes a similar move (p. 67). Here the claim is that people actually do have common or individual knowledge that they and/or others conform to some particular practice with 'might,' but they are unable to describe or express that knowledge in any way, even to themselves. As Lewis (1975) puts it, "[i]t is enough to be able to recognize conformity and non-conformity to [the] convention, and to be able to try to conform" (p. 25).

In order to be of any help, this ability to try to conform and to recognize conformity and non-conformity must be meant in a very light sense. If it required awareness of the nature of the practice to which one and/or others are conforming (or failing to conform), it wouldn't help at all with our examples, where that is just what is absent. As a result, what's required will have to be a mere ability to engage in the practice, and some kind of feeling (perhaps very nonspecific and subtle) that others' actions are usual or unusual. Then the idea would be that because English-speakers have a feeling that something is amiss when someone uses 'might' unconventionally, and because they are able to use it conventionally themselves, they have tacit common (or individual) knowledge that everyone (or just that individual) conforms to that practice.

One thing to note right away is that this move cannot possibly vindicate Lewis and Miller's *common* knowledge requirement. My subtle feelings of comfort or discomfort when others follow or violate a regularity and my ability to follow it myself seems like, at most, tacit knowledge that I and others conform, not tacit knowledge that others *know* about that conformity, and so on.

And if we applied this strategy to a requirement for individual rather than common knowledge, we'd be extending our concept of knowledge to the breaking point. If being able to do *A* and having a feeling that something is usual or unusual when others do or do not do *A* counted as having knowledge that I do (or everyone does) *A*, there does not seem to be any room for me to be able to

engage in *A* *without* having knowledge that I do (or everyone does) it. Engaging in *A* without that knowledge certainly can't mean that I'm unable to try to engage in *A*, so it must mean that I lack any feeling, however subtle, of familiarity or unfamiliarity when others do or do not do *A*. But it seems highly likely that for *any* practice in which I engage consistently, I will sometimes experience at least a very subtle and nonspecific feeling of familiarity or unfamiliarity when someone else does or does not engage in that practice. Even with a particularly subtle example, such as conventions about how to hold one's body while standing in line, people likely sometimes have a subtle feeling of ease or unease when someone engages or fails to engage in the practice to which they are accustomed. And yet, this seems like a clear case of a practice without knowledge of conformity. If Lewis's extension of the notion of knowledge were adopted, that case, too, would count as one in which we have knowledge of conformity, and it seems likely that no practices would be left in the category of those without knowledge of conformity.

In other words, this strategy would erase an antecedently plausible distinction between different kinds of practices: those in which we know that we and/or others conform, and those in which we do not. Note that I am not saying that this strategy erases the distinction between conventional and non-conventional practices. None of the accounts we're considering portray knowledge of conformity as the sole feature differentiating conventional practices from non-conventional ones. Rather, my concern is about a different distinction between kinds of practices: between practices in which we engage with the knowledge that we and others do so, and practices in which we engage without such knowledge. We'd be making our requirement for knowledge of conformity so minimal that no practice in which we actually engage would fail to satisfy it, and that in itself is a problem.

Leaving conventions aside for the moment, consider the rich vein of examples of practices without knowledge of conformity that sports provide. An experienced tennis player may swing their

racquet with a specific technique without knowing what it is they do. They are able to do it and they may feel a vague sense of familiarity if they encounter another player with the same swing, but they could not describe their own (or others') practice, or even identify the correct description if it were offered. On the other hand, another player may have considerable technical knowledge about the sport and perhaps about shoulder anatomy, and they may be able to accurately describe what they and others do when engaging in the same swing. The two tennis players engage in the same practice, the first without knowledge of doing so, and the second with such knowledge. It would be a mistake to paper over this straightforward distinction.

Moreover, just as with Lewis's appeal to potential knowledge, it would seem very odd to include a knowledge of conformity requirement in an account of conventions if 'knowledge' were meant in so very light a sense that knowledge of conformity was a feature of every practice in which we engage. Something stronger, again, is implied by the requirement's presence in the account.

Now of course, there are ways of cashing out the notion of tacit knowledge other than Lewis's proposal, but my contention is that any form of this strategy faces an insoluble dilemma: it will either require more of people than is present in cases such as the use of 'might' or 'but,' or it will require so little that it collapses the distinction between practices with and without knowledge of conformity.⁸ Thus, neither of Lewis's strategies can save any of the knowledge of conformity requirements we have considered.

The next step is to extend this section's criticism to apply to Gilbert's view. Recall Gilbert's (1989) requirement that members of a population in which a convention obtains "know that they jointly accept" the relevant principle (p. 373). English-speakers' inability to accurately describe what they do with 'might' and 'but' means that they cannot have knowledge of any particular principle that

⁸ For instance, Marmor (2009) exemplifies the second horn of the dilemma when he says that his 'davka' case (which is quite similar to our examples of semantic conventions) involves implicit "knowledge-how" that satisfies his requirement of knowledge of one's own conformity (pp. 67–70).

they have jointly accepted to guide their action with those words. If people can't identify a principle, they can't know that they all jointly accept it.

Thus, it seems as though the various knowledge of conformity requirements, and Gilbert's similar knowledge requirement, wrongly exclude some paradigmatic conventions, including semantic conventions that the authors in question ought to be particularly keen to accommodate.

4. Millikan's account of conventions

We need an account of conventions without a knowledge of conformity requirement. One might wonder whether one of the accounts of conventions already discussed could be suitable, if we removed the offending knowledge requirement. But unfortunately, the knowledge requirements in these accounts are either accompanied by other elements that are likely to cause similar problems, or too deeply entwined with the account to be excised.⁹ Instead, in this section I'll argue that Millikan's (2005, 2014) account of conventions, which does not require knowledge of conformity or anything similar, is on the right track, but complications stemming from her view that most linguistic conventions involve coordination and details of how she understands her key notion of reproduction prevent us from adopting her approach wholesale.

Millikan (2014) defines a convention as follows:

a pattern of behavior that is

- (1) handed down from one person, pair, or group of persons to others – the pattern is reproduced somehow – and
- (2) is such that if the pattern has a function, then it is not the only pattern that might have served that function about as well (p. 31).

⁹ Lewis's (1975) condition (2), along with the application of condition (6) to conditions (2)–(5), is likely to continue to wrongly exclude the examples of conventions we've discussed (and of course, as noted at the outset, there are other problems with Lewis's account). Similarly, Miller's (2001) key notions of joint action and collective ends make reference to additional mutual true beliefs that are likely to bring similar problems (pp. 59, 266). In Marmor's (2009) case, that conventions are rules requiring knowledge of one's own conformity seems like too deep and central a feature of his account to simply excise. Similarly, for Gilbert (1989), it is difficult to imagine how a group could issue a fiat to itself without members having some awareness that this has occurred.

Just from reading Millikan's account, it is probably clear that knowledge of conformity is not required. I'll say a bit more about each of her two requirements, to establish that no knowledge of conformity requirement sneaks in. But first, I want to note that I will refer to the bearers of conventionality as patterns of *activity* rather than behavior. Millikan (2005, 2014) seems to use 'behavior' and 'activity' interchangeably, but in my view, 'behavior' has connotations of complete observability that make it too narrow.

Millikan's condition (1) requires that a conventional pattern of activity be reproduced. According to Millikan, "[a] pattern has been *reproduced* if its form is derived from a previous item or items having, in certain respects the same form, such that had the model(s) been different in these respects the copy would have differed accordingly" (Millikan, 2005, p. 3). Millikan makes clear that reproduction is a causal notion (p. 31). If my activity is a reproduction of yours, then my activity has certain features because yours had them, in the causal sense of 'because' (*cf.* pp. 9, 55–56). Millikan emphasizes that reproduction can be unconscious (p. 5), and there's nothing else in condition (1) that would require knowledge of conformity. I could be causally influenced by someone else's activity without knowing exactly what activity I'm engaging in, and without knowing that anyone else engaged in that activity.

Millikan's condition (2) requires that conventions be arbitrary: if they serve any function at all, it must be one that at least one other pattern could have served approximately as well. The precise nature of conventions' arbitrariness is a major topic of its own, but for our purposes we just need to note that there's nothing in the claim that conventions must have alternatives that would require anyone to know that they or anyone else conform.

Millikan's account can accommodate the conventions that cause trouble for the other theorists we've discussed. For semantic conventions such as those involving 'of' or 'might,' Millikan could say that there is a pattern of usage that is reproduced (*i.e.*, people are causally influenced by others' past

usage), and yet people are not aware of exactly what it is they or others are doing. Moreover, the pattern they are reproducing has alternatives: there are many sounds aside from the one associated with 'might' that people could have used instead to serve the same function, as our planet's linguistic diversity amply demonstrates.

We might ask, at this stage, whether Millikan's account can satisfy the initial motivation that led Lewis to introduce his knowledge of conformity requirement, especially because that motivation centered on being able to accommodate linguistic conventions, which Millikan aims to do as well. Lewis thought that we needed to know which language others in our population use, in order to predict which language the individuals we now encounter will use. He thought this was necessary because of the diversity of possible languages. Millikan doesn't think we need to *know* which linguistic practice others use. Instead, we just need to be disposed to conform to what others have done in the past: we just need to be causally influenced to follow past precedent (Millikan, 2005, p. 7). This causal influence of past linguistic practice explains how we are able to communicate successfully despite the vast array of possible languages, and it certainly does not require prior agreement. So, Millikan's account satisfies the heart of Lewis's motivation, but without excluding the key class of conventions his account excludes.

This is very promising, but there is another feature that Millikan thinks most linguistic conventions share that prevents us from adopting her approach wholesale. Millikan (2005) says that along with satisfying her basic account of conventions, most linguistic conventions have an additional feature: they involve coordination. The only examples she offers of non-coordinative linguistic conventions are along the lines of expletives and interjections, suggesting that on her view, it really is nearly all linguistic conventions that are coordinative (p. 2). This is an area of partial, limited agreement with Lewis.

Because Lewis's (1969/2002) definition of coordination requires a highly complex "system of concordant mutual expectations," it might seem that Millikan's view that most linguistic conventions involve coordination would lead to problems analogous to those facing knowledge of conformity requirements (p. 33). But, Millikan (2005) offers a much simpler definition of coordination:

There is a need for coordination when:

1. members of a group ... have a purpose in common;
2. achieving this purpose requires actions by each...;
3. more than one combination of actions will achieve the purpose;
4. the set of workable combinations fails fully to determine what any single partner's contribution must be independently of the actions performed by the others.

Coordination is achieved if the ... combined actions achieve the common purpose (p. 9).

Here, coordination is just a matter of being in a situation with a particular structure, in which the shared purpose is actually achieved. As a result, "[t]here is no need for the various parties ... even to recognize the problem as a coordination problem, let alone to think about one another's thoughts" (p. 57).¹⁰

Despite this less demanding definition of coordination, there are problems with Millikan's claim that most linguistic conventions involve coordination, as we can see by considering one of her own examples: the indicative mood. Recall that condition (1) of Millikan's definition of coordination requires that the people involved have "a purpose in common." She tells us that the purpose in common that a speaker and hearer coordinate to achieve when a speaker uses the indicative mood to express some proposition p is "that [the hearer] should become informed about p " (Millikan, 2005, p. 45). I'll grant, for now, that the speaker's purpose in expressing that p in the indicative mood is often that the hearer become informed about p , but this cannot possibly be an accurate characterization of the hearer's purpose. Only in a very small percentage of cases could we plausibly describe a hearer as approaching the speaker's utterance with the purpose of becoming informed about the specific proposition that the speaker will assert. Unless the conversation is highly constrained, the exact

¹⁰ See Guala, 2020 for an argument that coordination actually works better, in at least some circumstances, when people do not think about each other's thoughts.

proposition that the speaker expresses will be a surprise to the hearer, even if the topic of conversation is not.

In response to this worry, Millikan could say that the shared purpose is to transfer information from speaker to hearer (as she says about “the indicative-mood form itself” (Millikan, 2005, p. 45)), but then we have a more accurate characterization of the hearer’s purpose at the expense of the speaker’s. The hearer’s purpose may often be to get some information or other transferred to her from the speaker, or perhaps to get information on a particular topic. But the speaker’s purpose isn’t just to transfer any old information whatsoever, or even to transfer some information or other on a given topic; rather, it’s to transfer the information that *p*.

Moreover, even if we grant that the hearer’s purpose in approaching the utterance is to become informed about a specific proposition, problems remain with construing the speaker and hearer’s shared purpose as the hearer becoming informed about *p*. The speaker’s purpose in making the utterance will typically be for the hearer to become informed about *p* in the sense of ‘learning that *p* is true,’ whereas the hearer’s purpose will typically be to become informed about *p* in the sense of ‘learning whether *p* is true.’ It may seem harmless to say that learning that *p* is true is just one particular way of learning whether *p* is true, and thus that the interlocutors share the purpose of the hearer learning whether *p* is true. But again, this is an inaccurate characterization of speakers’ typical purposes. If I tell you that I’m happy and you learn from my utterance that I’m not happy (perhaps because you know that I only feel moved to proclaim my happiness when I’m in denial about how unhappy I am), my purpose as a speaker will have been thwarted, not satisfied. When a speaker tells a hearer that *p*, their purpose in doing so is for the hearer to learn that *p* is true, not just to learn whether *p* is true. A hearer’s purpose, on the other hand, is just to learn whether *p* is true. They are not coordinating.¹¹

¹¹ Schiffer (1972) and Miller (2001) raise concerns about coordination similar to the ones in the preceding paragraphs, though not in relation to a less intellectualized understanding of coordination like Millikan’s (pp. 151–153; p. 70).

Similar problems arise with respect to word-level semantic conventions. We might plausibly characterize a speaker's purpose with 'tree' as 'getting the hearer to think of trees.' But again, unless the conversation is highly constrained, many of the specific objects and properties the speaker designates will be a surprise to the hearer, which makes it implausible that the hearer typically approaches an utterance of 'tree' with the antecedent purpose of being made to think of trees. The only purpose the interlocutors can be said to share is something very general, such as information being transferred to the hearer. And that again seems misleading on the speaker's side. The speaker's purpose in uttering 'tree' is not just that communication occur, but rather, her purpose is to communicate something *about trees*. The point generalizes straightforwardly to a vast range of other words, undermining Millikan's claim that most linguistic conventions involve coordination.

One might object at this stage that even if the purpose actually shared by speaker and hearer is overly general for the speaker, it is nonetheless a purpose the interlocutors share, and thus, such conventions do involve coordination. But notice how this move would lump together all semantic conventions. What's the purpose of the convention for 'tree'? For some information or other to be transferred from the speaker to the hearer. What's the purpose of the convention for 'stapler'? For some information or other to be transferred from the speaker to the hearer. We're then left without an explanation of why languages have so many different words, if they all have the same purpose.

Here, the objector might respond by saying that although all of our word-level semantic conventions have the same coordinative purpose, we need many different ones because each makes a different contribution to achieving that purpose, such as communicating something about trees, in particular. But then, that unique contribution to communication would be each convention's primary purpose (*i.e.*, it is what that convention, in particular, does), and that purpose is not coordinative.

So, Millikan is wrong to think that most linguistic conventions are coordinative. This conclusion might not seem especially troubling, given that coordination is not actually built into her

account of conventions. However, recognizing the widespread non-coordinative nature of linguistic conventions leads to a significant change in how we can apply the central concept in Millikan's account of conventions, reproduction, to linguistic conventions.

Millikan emphasizes a particular kind of reproduction, counterpart reproduction, as the mechanism by which linguistic conventions spread. Ballroom dancing is a paradigmatic example. One partner who already knows a dance might spread it to another who does not, by leading the dance while the other follows (Millikan, 2005, p. 4). For linguistic conventions, the speaker's role in the "dance" is to produce an utterance, and the hearer's role (for an indicative utterance) is to interpret the speaker and form a corresponding belief (p. 6).

If we look closely, it becomes evident that only coordinative conventions can spread by counterpart reproduction. Millikan (2005) describes counterpart reproduction as involving a single conventional pattern, of which multiple individuals perform a part (p. 4). Thus, *pairs* of dancers or interlocutors copy past *pairs* of dancers or interlocutors to complete the pattern. To satisfy Millikan's definition of convention, that overall multi-person pattern would need to have *one* function (if any) that it is reproduced to achieve. In other words, the participants must have a "purpose in common" that they achieve through their combined actions. Counterpart reproduction thus implies coordination.

So, the considerations that pushed us away from coordination also push us away from Millikan's preferred form of reproduction for linguistic conventions. Happily, we can consider another kind of reproduction that applies to non-coordinative conventions: direct copying, where someone is just straightforwardly influenced by someone else's activity to engage in activity with the same feature(s) (Millikan, 2005, p. 4). A key difference between direct copying and counterpart reproduction is that in direct copying, the copying does not occur by means of one individual initiating part of a multi-participant pattern and then influencing someone else to fit in to complete it. Rather, in direct

copying, the copied individual's activity just causes the copier to engage in activity with the same feature(s) on her own. Crucially, there is no need for a shared purpose.

Shifting from counterpart reproduction to direct copying is far from straightforward, because Millikan identifies two advantages of counterpart reproduction that she seems to think direct copying lacks. First, Millikan (2005) argues that counterpart reproduction “results in standardization of forms, more easily than ... direct copying.” The idea is that if I learn to waltz by dancing with someone who already knows the other half of the pattern, I will be constrained in how much I can deviate from the pattern by the need to conform pretty closely to be able to continue dancing with my partner at all. On the other hand, if I just directly copy another individual's solo dance moves, my version can diverge from what I'm copying without any consequences for my ability to complete my activity. As a result, over time, the solo dance move will experience quite a bit more “drift” away from its original form than a partnered dance move would, and thus is likely to be much less standardized (pp. 4–5). This matters for our purposes because semantic conventions are, of course, quite standardized. Complete strangers are often able to communicate, and vast numbers of people can understand the same books, podcasts, and films. If switching from counterpart reproduction to direct copying undermined our ability to explain this standardization, that would be disadvantageous.

Second, Millikan argues that counterpart reproduction can very clearly be unconscious. We become “*adapted* to an environment” in which others produce their halves of conventional patterns, and we develop the habit of performing the other half when someone else initiates (Millikan, 2005, pp. 6, 10, 57). Other people's activity molds us into playing our part, over time. This feature allows Millikan's account to accommodate conventions without knowledge of conformity. If direct copying required awareness of the pattern one is reproducing, then it would reinstate a knowledge of conformity requirement. I'll address these two concerns in turn.

First, standardization. I'll begin by noting that it is important not to exaggerate the degree of standardization in natural languages. There are regional differences in the use of many words, sometimes generating misunderstanding among people who ostensibly speak the same language. And most language-users have habitual linguistic idiosyncrasies. But yes, there is quite a bit of standardization in contemporary, global natural languages, more than what we would expect from simple copying over such a large span of time and across so many individuals. This standardization can be explained by people's use of a variety of mechanisms to increase the accuracy of their copying. Dictionaries and formal schooling both aid individuals in making their copies of others' language use more accurate. Public media widely disseminate the language use of a small subset of the population, amplifying the causal influence of those individuals and thus helping to slow drift. These mechanisms promote standardization, without using counterpart reproduction. Once we take them into account, there is no explanatory gap, in my view.

Next, we need to consider whether direct copying can be unconscious. Even though Millikan's general definition of reproduction does not require knowledge of conformity or of anything else, there is still always an open question about whether a particular kind of reproduction might. With respect to counterpart reproduction, Millikan seems to think we learn the pattern as a part of our general ability to get by in the world, by learning to fit in with others who initiate the pattern. This is part of what I take Millikan (2005) to mean when she says that "[r]eproduced patterns are often learned rather in the manner of skills" (p. 5). They are not skills because they are not in general the best way of achieving their function, but learning them can be just as unconscious as picking up a necessary skill because the standardization in others' activity makes them the best way to achieve their function within a given group.

So, the question is: can and do we unconsciously imitate others' activity directly, just as we unconsciously modify our activity to follow someone's lead in a ballroom dance? Consider your

experiences with a longtime friend who has a relatively new, but serious, romantic partner. Your friend's mannerisms might subtly change: they might smile a bit differently, or brush their hair aside in a new way. Then, you might notice that the romantic partner shares those gestures. Your friend has picked up some of their partner's mannerisms, almost certainly without intention or awareness. And this is not a matter of coordination; there is no shared end being achieved. Rather, one person's activity is just directly, unconsciously influencing the other's. We influence each other in this way all the time. For more examples, consider the way mannerisms and other subtle habits spread among other people who associate with each other, whether categorized geographically, institutionally (such as coworkers, students at the same school, or members of a religious organization), or culturally.

I'll also note that direct copying is a particularly plausible mechanism for linguistic conventions because like many of our linguistic abilities, our level of facility with it sets us apart from other primates. According to empirical work helpfully summarized by Richard Moore (2013), chimpanzees are adept at observing someone else achieving a goal and then trying to achieve the same goal, but they typically invent their own means to the goal rather than copying the means used by whomever they observed. Only rarely do they directly copy another individual's specific action. Humans, on the other hand, frequently copy others' means as well as others' goals (pp. 497–8).¹² The fact that our ability to engage in this kind of copying sets us apart from other primates recommends it as an explanation for our equally distinctive communicative practices. In other words, part of the reason why humans have complex natural languages whereas other primates do not is that humans are adept at directly copying others' activity.¹³

¹² Millikan (2005) makes a similar point (pp. 56–57). Moore (2013) suggests that chimpanzees must be thinking to themselves “[Effect] *E* can be achieved in [situation] *S*,” whereas humans think, “Performing *A* in *S* is a way of doing *E*,” implying that this kind of copying cannot be unconscious (p. 498). However, I see no reason to assume that humans or chimpanzees must think anything in particular while imitating others, especially given the examples of unconscious direct copying provided above.

¹³ Although our propensity to copy others' actions sets us apart from other primates, it does not set us apart from all other animals. For example, there is evidence that prairie dogs may have copied, conventional languages (Devitt, 2021, pp. 33–34). I do not mention this phenomenon in the main text because prairie dogs are not close evolutionary relatives of ours,

So, it looks as though shifting away from coordination, and thus from counterpart reproduction to direct copying, will not deprive us of any of the advantages Millikan attributes to counterpart reproduction. However, we cannot simply adopt Millikan's notion of direct copying, due to an additional problem with her notion of reproduction.

As a reminder, Millikan (2005) defines reproduction (of which direct copying is a species) as follows:

A pattern has been *reproduced* if its form is derived from a previous item or items having, in certain respects the same form, such that **had the model(s) been different in these respects the copy would have differed accordingly** (p. 3).

The bolded, counterfactual element of this definition is too demanding. Imagine that I directly copy someone else's choice of what to serve guests for dinner. Millikan's definition of reproduction requires that had the cook served something else, I would still have copied her. This might sometimes be the case, such as when the cook has a high status in my social group. But it also might not be the case: I might have an affinity toward someone's particular menu selection and copy her, and still not be such that I would have copied *whatever* meal she served. In my view, copying requires only the right kind of causal connection between one person's activity and another's; how the second person would have acted if the first had acted differently just isn't relevant (*cf.* Stotts, 2021, p. 30).

Millikan's approach to conventions is clearly a major improvement over accounts requiring knowledge of conformity.¹⁴ But we saw reason to move away from her view that most linguistic conventions involve coordination, which then revealed that the kind of reproduction that gives rise

and so their inability to develop languages as complex as ours, even though they are able to copy, is not particularly mysterious. What does call out for explanation is the vast difference in the communicative abilities of humans among primates, given how similar we otherwise are. And humans' ability to copy seems to be part of that explanation.

¹⁴ I should note that Davis (2003) offers an account of conventions that allows for some conventions without knowledge of conformity. He provides a list of mechanisms by which conventions are perpetuated, one of which (enculturation) hints at something akin to Millikan's notion of reproduction, though without the depth in which she develops that notion and without portraying it as a unifying factor in all conventions. The rest of his mechanisms follow Lewis in presupposing or explicitly requiring knowledge of conformity (pp. 207–211).

to such conventions is not counterpart reproduction, but direct copying. And we discussed a further problem with Millikan's definition of reproduction that applies to direct copying, too.

5. A revised account of conventions

Drawing from the discussion in the previous section, I'll now offer a revised definition of copying and then a revised account of conventions, highlighting several advantages of the resulting account over Millikan's.

First, I want to note an aspect of Millikan's definition of reproduction that is crucial to preserve. Millikan's definition mentions "certain respects" in which a copy's form matches the original. As a result, the causal relationship we're tracking when talking about reproduction obtains between *features* of the original and *features* of the copy. This matters because, as Millikan (2005) puts it, "[a] reproduction is always a reproduction only in certain respects," and thus a single instance of activity can be a copy of multiple different past instances of activity, in different respects (pp. 30–31).

With this in mind, I propose the following preliminary definition of copying (I prefer 'copying' to 'reproduction' because I think the biological analogy easily becomes overextended): activity *Y* is a copy of activity *X* if and only if one or more features of *X* caused the same feature(s) in *Y*. Here I appeal only to an actual causal relationship and exclude the counterfactual element that caused problems for Millikan.

The 'same feature(s)' part of the definition is intended to be undemanding, allowing poor copies to count as copies. For instance, if I radically mispronounce a word, the only feature it shares with the past utterances from which it was copied may be something like "being a sound made by a mouth that is audible to humans." This would still be enough for it to count as a copy of those past utterances. The notion of cause is similarly undemanding. There is no minimum amount of influence the past activity must have exerted on the present activity in order for the present activity to count as

a copy of it. This means that copying admits of degrees: my activity might be a copy of two different dancers' moves, but of one to a much greater degree than the other, due to a higher number of shared features or a greater degree of causal influence.

The definition just offered is preliminary because we still need to differentiate copying from cases in which a species has evolved to engage in activity with certain features. For instance, imagine that a member of a certain species of cat lunges at a certain speed because lunging at that speed allowed past members of the species to get food and survive long enough to sexually reproduce, thereby passing on genes that cause the present cat to lunge at the same speed. In that case, my preliminary definition would imply that the present cat is copying those past cats (with each past cat exerting only a tiny amount of influence). But this doesn't seem like the kind of copying that gives rise to conventions.

Millikan (2005) herself avoids this issue, but her way of doing so relies on the problematic counterfactual component of her account and is thus not available to us. She would say that although there is a causal connection between past cats' behavior and the present cat's behavior, if any one of those past cats had lunged at a different speed due to an injury or illness, that would not have caused the present cat to lunge differently too, because the genetic information that leads to the inherited lunging behavior is the only thing those past cats could pass on. So, this case does not fit Millikan's definition of reproduction (p. 30). But without the counterfactual component of Millikan's definition, we just have a requirement that there be shared features, with a causal connection. And that does obtain, albeit indirectly, between past cats' behavior and the behavior of our present cat.

We need to add something to our definition of copying, but something weaker than Millikan's counterfactual component. My proposal is to require that when copying occurs, the mechanism of causation must do its work within the scope of the copying individual's own life. The causal influence of past cats' lunging on the present cat's lunging was carried out via the mechanism of genetic

inheritance. That causal mechanism's work was complete before the present cat was born, so it's not an instance of behavior being copied in my sense. But when I see someone else dancing and copy her moves, or I am exposed to a certain way of using 'of' and begin to use it the same way myself, the mechanism of causal influence operates within the scope of my own life.

The virtue of this way of handling the matter, in my view, is that it is very direct. But it may seem to have the vice of being *ad hoc*. To the contrary, I think that it captures something crucial and independently significant about the nature of copying: when we talk about activity being copied, we're talking about one *individual* copying some other individual. So, we want the copying to be done *by* the copying individual, as the kind of being that they are, with all their innate features in place. Before the lunging cat had the genes that compel it to lunge in the particular way it does, there wasn't any cat *there* to do any copying. The cat just emerged as a being that was going to lunge in that way. *It* didn't copy any other cats' lunging behavior, though some genetic copying (which is a different form of copying altogether) occurred in the process of the cat's conception.

My final definition of copying, then, is the following: activity *Y* is a copy of activity *X* if and only if one or more features of *X* caused the same feature(s) in *Y*, where the mechanism of causation operates within the scope of the copying individual's own life. This definition preserves the virtue of Millikan's definition on which we've been focusing: copying can still be unconscious. As on Millikan's definition, features of past activity can cause my present activity to have the same features, without my realizing that this has occurred or being aware of the nature of my activity. And of course, this definition eschews Millikan's overly demanding counterfactual requirement.

Relying on this revised definition of copying, I'll now present my overall account of conventions.

Doing A is a convention within a group in virtue of the following conditions being satisfied:

- (1) Within the group there is widespread, interconnected copying of doing A as a way of doing something further, B .
- (2) There is at least one other way of doing B that is at least approximately as good and at least approximately as accessible to the group as doing A , independent of the dominance doing A has gained due to being copied.¹⁵

Although some conventions may involve coordination, I don't take that to be a feature of any particular importance with respect to their nature as conventions; it's just one way of describing the manner in which certain conventions satisfy condition (2).

Despite owing much to Millikan, my account of conventions differs from hers in some significant ways, beyond the change to the definition of copying. As a reminder, Millikan (2014) tells us that a convention is:

a pattern of behavior that is

- (1) handed down from one person, pair, or group of persons to others – the pattern is reproduced somehow – and
- (2) is such that if the pattern has a function, then it is not the only pattern that might have served that function about as well (p. 31).

I'll discuss some additional differences between my account and hers, highlighting what I take to be the advantages of my approach.

One notable difference between my account and Millikan's is my talk of "doing A as a way of doing something further, B ." This replaces Millikan's talk of the function a convention serves. Something along these lines is needed because in order to capture the idea of equally good alternatives to some practice, we need to know: equally good with respect to accomplishing *what*? But on the kind of etiological approach to functions to which I am sympathetic, a function is an *effect* (Wright, 1973,

¹⁵ Condition (2) is modeled on my Alternatives Condition for Conventions (Stotts, 2017, p. 880).

pp. 160–161).¹⁶ Shifting toward the ‘as a way of’ language and away from the ‘function’ language is advantageous because it allows the relationship to sometimes be constitutive rather than causal. Consider waving: it’s a way of greeting, but greeting isn’t some further effect.¹⁷

A related difference is that Millikan allows that some conventions may lack functions altogether, whereas my definition requires that all conventions are done as a way of doing something further. This change is an improvement because, as Millikan’s account stands, any activity that we copy others in doing as an end in itself counts as conventional. For instance, we copy each other in seeking pleasure. When I see others seeking pleasure, I am often motivated to seek some for myself, and not because of any further function that my pleasure serves. Seeking pleasure satisfies Millikan’s account of conventions, but is not conventional, because it lacks conventions’ characteristic arbitrariness. Moreover, I find Millikan’s purported examples of conventions without functions unconvincing. Millikan offers the example of using red and green to decorate for Christmas, but this clearly seems to be done as a way of making it publicly recognizable that one is celebrating Christmas. Similarly, her example of smoking cigars when a boy is born clearly is done as a way of demonstrating excitement about having a baby boy (Millikan, 2005, p. 7).

Despite these advantages, there is a concern that might arise about the notion of “doing A as a way of doing B .” One might doubt whether there will be single, determinate values for A and B across the entire group, for certain semantic conventions. My view that each instance of copying only needs to preserve one feature of what was copied allows for variation within a group, even within a relatively short span of time. There could be subtle variation in how people in the same group pronounce a given word, which would be variation with respect to A . And the meaning of the word—

¹⁶ Millikan’s (1989) own account of functions pushes in a similar direction (p. 288).

¹⁷ One might wonder whether the ‘as a way of’ talk inadvertently reintroduces a knowledge of conformity requirement. Can I do something as a way of doing something else without being aware of what I’m doing? For a response to this worry, see Stotts, 2021. I’ll also note that my use of the ‘as a way of’ notion is indebted to Israel *et al.*, 1993.

what group members use the word as a way of expressing—could vary subtly from person to person in a similar way, which would be variation with respect to *B*. We can appeal to the standardization-promoting mechanisms discussed in Section 4, but nonetheless, there is room for subtle variation. How, then, could there be a determinate value, at the group level, for either *A* or *B*?

With respect to the values for *A*, I am content to follow Millikan (2005) in seeing token instances of activity as sorted into types by causal history (pp. 33–35, 60–62). So, two instances of activity count as tokens of the same type precisely because their causal histories of copying overlap. This allows for subtle variation in word pronunciation across a group. As long as the causal connections of copying are preserved, the token activities are of the same type.

For the values for *B*, I take a different approach. The value of *B* will always be the largest shared element of what the majority of individuals who do *A*, do it as a way of doing. For example, each individual in a particular group might utter ‘blue’ as a way of designating slightly different parts of the color spectrum. The meaning of ‘blue’ in the group will be the largest part of the color spectrum that most of the individuals’ use has in common.¹⁸

Another key difference between my account and Millikan’s, for which I argue elsewhere, is that I require that a convention’s alternatives must be not just equally good, but also equally *accessible*, where an alternative’s accessibility to an individual or group has to do with “the degree of ease with which that individual or group can encounter and use it” (Stotts, 2017, p. 875). This qualification ensures that only practices with *genuine* alternatives—ones that participants realistically could have used instead—count as conventional.

¹⁸ This way of dealing with the matter pushes me in the direction of semantic minimalism (Borg, 2004, 2012; Cappelen and Lepore, 2005). Developing that connection would take us too far afield, but I do want to acknowledge it. I also want to note that in addition to the subtle variation within a group at a single (span of) time that we have been discussing, similar variation can happen across time. The causal history strategy for values for *A* works equally well for variation across time. On the other hand, once the value for *B* at the group level has changed with the passage of time, my view is that the convention would then be changed. This is what happens when a word’s meaning shifts over time.

The preceding divergence from Millikan gives rise to another. Because accessibility must always be accessibility *to* some individual or group, I index conventions to groups. Millikan (2005) does not index conventions to groups, and in fact, she criticizes Lewis's indexation of conventions to groups (or as he would say, populations) (p. 29). But her criticism applies specifically to Lewis's requirement of a *regularity* within a population, which is not part of my account. I require just that a convention be *widespread* within a group, a much lighter requirement. Because the idea that conventions belong to groups or populations is relatively uncontroversial among other authors, I won't say more about it here.

I do want to say more about the just-mentioned requirement that the copying that gives rise to a convention be *widespread*, and about my related requirement that the copying be *interconnected* within the group. Millikan does not specify how many times a pattern must be copied in order to be conventional, leaving her open to the objection that she counts something that is copied only once as a very short-lived convention. In response, Millikan (2005) could note that conventionality admits of degrees, and so a pattern copied only once would be only very slightly conventional (*cf.* p. 49). But still, while I agree that conventionality admits of degrees along numerous dimensions, I also think there must be a copying threshold below which a convention does not yet obtain. This threshold should be vague, given that whether some practice is established enough to be conventional in a given group is itself vague (*cf.* Davis, 2003, pp. 218–219). My proposal is to require that the copying be widespread, where 'widespread' means that the majority of group members have encountered some instance of the practice being copied within the group, even if they have never personally copied it. The copying must also be interconnected to avoid counting as conventional cases in which many people in some group engage in the same activity for independent reasons, and then each of them is copied only once by someone else. This would seem like just a coincidence, rather than a convention (*cf.* Stotts, 2021, p. 30).

A final advantage of my account over Millikan's lies in my qualification that a convention's alternatives must be equally good and equally accessible only *independent of the dominance the convention gained due to being copied*. This qualification is needed to allow that once a convention is well established, it may *then* become the best and most accessible way of doing something in a particular group (*cf.* Stotts, 2017, p. 879). Consider semantic conventions: now that certain conventions with 'tree' are well established, that word *is* the best and most accessible way to achieve its conventional purpose, in certain groups. Millikan (2005) makes remarks that suggest she would be amenable to such a qualification, but I think it is worth making explicit (pp. 6, 57).

The final step is to make the case that my account can still accommodate conventions without knowledge of conformity. To make the task more concrete, again consider 'might.' There is widespread, interconnected copying within various groups of using 'might' in a particular way, although the participants lack awareness of the precise nature of their activity. This lack of awareness is (as we saw above) perfectly compatible with my revised definition of copying. And, there are alternatives in the right sense: there are other sounds people in those groups could have used instead that were initially equally good and equally accessible. The existence of such alternatives in no way requires anyone to know that they or others conform. And, to consider a non-linguistic example, there is widespread, interconnected copying within different groups of different practices about how often to smile at others, again without awareness. And again, other smiling practices were initially equally good and accessible. So, it seems we have arrived at an account of conventions that is stronger overall than Millikan's, and which still has the advantage over Lewis and others of accommodating conventions without knowledge of conformity.

I want to end with a caveat: it is easy for the de-intellectualized approach to conventions I promote to be interpreted as a de-valuing of the intellectual side of human sociality, but that is not my intention. Our inferences about what others will do in the future, our sense of mutual awareness in

our groups that certain practices are *our* practices, and many of the other mental phenomena that Lewis and others emphasize actually are important parts of the broader social picture. My view is that, legitimate as these elements of sociality are, they are not part of one particular social phenomenon: conventions.

References

- Borg, E. (2004). *Minimal semantics*. Oxford University Press.
- Borg, E. (2012). *Pursuing meaning*. Oxford University Press.
- Burge, T. (1975). On knowledge and convention. *The Philosophical Review*, 84(2), 249-255.
- Cappelen, H., & Lepore, E. (2005). *Insensitive Semantics*. Blackwell.
- Davis, W. A. (2003). *Meaning, expression and thought*. Cambridge University Press.
- Devitt, M. (2021). *Overlooking conventions: The trouble with linguistic pragmatism*. Springer Nature.
- Gilbert, M. (1989). *On social facts*. Routledge.
- Guala, F. (2020). Solving the hi-lo paradox: Equilibria, beliefs, and coordination. In *Minimal cooperation and shared agency* (pp. 149-168). Springer Nature.
- Hayduk, L. A. (1983). Personal space: Where we now stand. *Psychological Bulletin*, 94(2), 293–335.
- Israel, D., Perry, J., & Tutiya, S. (1993). Executions, motivations, and accomplishments. *The Philosophical Review*, 102(4), 515-540.
- Lewis, D. (1969/2002). *Convention: A philosophical study*. Blackwell.
- Lewis, D. (1975). Languages and language. In K. Gunderson (Ed.), *Language, Mind and Knowledge: Minnesota Studies in the Philosophy of Science, Vol 7* (pp. 3-35). University of Minnesota Press.
- Marmor, A. (2009). *Social conventions: From language to law*. Princeton University Press.
- Miller, S. (2001). *Social action: A teleological account*. Cambridge University Press.
- Millikan, R. G. (1989). In defense of proper functions. *Philosophy of Science*, 56(2), 288–302.
- Millikan, R. G. (2005). *Language: A biological model*. Oxford University Press.
- Millikan, R. G. (2014). Deflating socially constructed objects: What thoughts do to the world. In M. Gallotti & J. Michael (Eds.), *Perspectives on social ontology and social cognition* (pp. 27-39). Springer.
- Moore, R. (2013). Imitation and conventional communication. *Biology & Philosophy*, 28(3), 481-500.
- Schiffer, S. (1972). *Meaning*. Oxford University Press.
- Sorokowska, A. et al. (2017). Preferred interpersonal distances: A global comparison. *Journal of Cross-Cultural Psychology*, 48(4), 577–592.

Stotts, M. H. (2017). Walking the tightrope: Unrecognized conventions and arbitrariness. *Inquiry*, 60(8), 867-887.

Stotts, M. H. (2021). Behavioral foundations for expression meaning. *Topoi*, 40(1), 27-42.

Wright, L. (1973). Functions. *The Philosophical Review*, 82(2), 139-168.