

Walking the Tightrope: Unrecognized Conventions and Arbitrariness

Megan Henricks Stotts

This is an Accepted Manuscript of an article published by Taylor & Francis in *Inquiry* on March 11, 2016, available online: <https://doi.org/10.1080/0020174X.2016.1140072>.

Abstract:

Unrecognized conventions—practices that are conventional even though their participants do not recognize them as such—play central roles in shaping our lives. They range from the indispensable (e.g. unrecognized linguistic conventions) to the insidious (e.g. some of our gender conventions). Unrecognized conventions pose a challenge for accounts of conventions because it is difficult to incorporate the distinctive arbitrariness of conventions—the fact that conventions always have alternatives—without accidentally excluding many unrecognized conventions. I develop an Accessibility Requirement that allows us to account for both arbitrariness and unrecognized conventions. Specifically, I argue that a conventional practice must have at least one alternative that is at least approximately as good and at least approximately as accessible as the conventional practice itself, independent of the dominance the practice gained as it became conventional. In the course of arguing for this requirement, I also show that two prominent accounts of conventions, David Lewis’s and Ruth Garrett Millikan’s, run into problems with capturing the arbitrariness of conventions. The Accessibility Requirement opens the door to improved accounts of conventions by precisely identifying the way in which conventions are arbitrary.

I. Introduction

Conventions are a subtle but powerful part of social reality: they can shape our lives without our realizing that they are mere conventions. Traditional gender practices were once widely believed to be prescribed by nature, and the realization that they had actually just been conventional all along was groundbreaking. For example, wives completing household tasks while husbands worked outside the home was considered the only acceptable way of dividing up labor within a family, and the realization that this was merely a conventional practice was not an easy one.

Linguistic conventions are also striking. Most language-users are aware that the meanings of words are conventional, but many of us are unaware of less obvious linguistic conventions. Consider the convention of using rising intonation when asking a question. Speakers of many languages do this without realizing that it is merely conventional and that some languages do not use intonation in this way. Or consider the English convention of starting a syllable with no more than three consecutive consonants. Most English-speakers probably don’t realize that they engage in this practice, let alone that it is conventional.

I will use the term ‘unrecognized conventions’ for practices that are conventional even though most of their practitioners do not recognize them as such. Any account of conventions must accommodate the many unrecognized conventions that shape our lives, whether this is beneficial (as in the case of language) or harmful (as in the case of certain social restrictions).

An account of conventions must also capture their *arbitrariness*: the fact that conventions have alternatives.¹ For example, the practice of planting grass around houses is conventional partly because there are plenty of other ways we could fill that space instead. On the other hand, the non-

¹David Lewis, Ruth Garrett Millikan, Tyler Burge, and Andrei Marmor share the starting assumption that conventions are arbitrary. See Lewis, *Convention*, 69–70; Millikan, *Language: A Biological Model*, 7; Burge, ‘On Knowledge and Convention’, 249; and Marmor, ‘On Convention’, 349. On the other hand, Seumas Miller and Margaret Gilbert deny that conventions have to be arbitrary. See Miller, *Social Action*, 115 and Gilbert, *On Social Facts*, 340–1.

conventional practice of putting roofs on houses is *not* arbitrary: a roof is the best way to provide adequate shelter from inclement weather.

Unfortunately, attempts to capture the arbitrariness of conventions tend to exclude many unrecognized conventions. First, simply saying that conventions must have alternatives creates a serious risk of *false positives*. False positives are non-conventional practices that get erroneously classified as conventional. Then, as we'll see, this risk of false positives creates pressure to require that participants be *aware* of their convention's alternatives. But if we require awareness of alternatives, we exclude many unrecognized conventions, creating equally problematic *false negatives*. False negatives are conventional practices that get erroneously classified as non-conventional. I aim to solve this problem by developing an Accessibility Requirement: a convention's alternatives must be at least as accessible as the conventional practice itself, in a very particular sense of 'accessible'.

In Section II, we'll discuss the difficulties that arise when we try to capture the arbitrariness of conventions, showing how attempts to avoid the risk of false positives send us careening toward the troubling false negatives. Here we'll draw connections to work by David Lewis and Ruth Garrett Millikan, who provide two of the major approaches to conventions.² The primary difference between their views is that Lewis sees conventions as necessarily rational, whereas Millikan sees them as arising from our non-rational copying behavior. It is instructive to discuss their views side by side because both authors aim to include unrecognized conventions³ but nonetheless risk excluding them.

We might think of the project of incorporating arbitrariness into an account of conventions as walking a tightrope. Leaning too far in either direction, toward the false positives or false negatives, amounts to falling off. Section III introduces, develops, and applies the Accessibility Requirement—the metaphorical balance pole that allows us to safely walk the tightrope. Then Section IV brings more detail about why neither Lewis nor Millikan is able to walk the tightrope.

Tyler Burge argues that Lewis's analysis of convention cannot accommodate unrecognized conventions.⁴ Burge also briefly proposes a way of capturing the arbitrariness of conventions without excluding unrecognized conventions. We will be building on Burge's insights, but as we'll see, quite a bit of philosophical work still needs to be done to ensure that we properly categorize all of the crucial cases and arrive at a complete understanding of the arbitrariness of conventions. The Accessibility Requirement provides this needed development. We'll also see that Millikan's account of conventions, which in many ways takes to heart Burge's criticism of Lewis, nonetheless is incompatible with a central aspect of the Accessibility Requirement. The Accessibility Requirement will open the door to better accounts of conventions by identifying the precise sense in which conventions are arbitrary.

²Lewis's account of conventions has been the most influential, with many other theorists criticizing but ultimately building upon his work. See, for instance, Schiffer, *Meaning*; Bennett, *Linguistic Behaviour*; Miller, 'Rationalising Conventions'; Miller, *Social Action*; Marmor, 'On Convention'; Marmor, *Social Conventions*; Sillari, 'Rule-Following as Coordination'; Tummolini et al., 'A Convention or (Tacit) Agreement'; and Guala, 'Normativity of Lewis Conventions'. Millikan's account of conventions has been less influential thus far, but it provides an interesting alternative to Lewis's approach and has received some discussion. See Chomsky, 'Reply to Millikan'; Epstein, 'Review of Ruth Millikan'; Tomasello, 'Conventions Are Shared'; Cain, 'Conventions and Their Role'; Moore, 'Imitation and Conventional Communication'; Azzouni, 'Defense of Logical Conventionalism'; and Witek, 'Interactional Account'. Another importantly different approach to conventions, which we won't discuss here, is Margaret Gilbert's in *On Social Facts*.

³Lewis, *Convention*, 63–4; Lewis, 'Languages and Language', 25; Millikan, *Language: A Biological Model*, 148; and Millikan, 'Deflating Socially Constructed Objects', 34.

⁴Burge, 'On Knowledge and Convention'.

II. The Tightrope: False Positives and False Negatives

In this section we will come to understand the difficulty of avoiding both false negatives and false positives when attempting to capture the arbitrariness of conventions. But first, a clarification about the false positives. Because arbitrariness is just a necessary condition for conventionality, we would need a complete account of conventions (which is beyond the scope of the present discussion) to definitively show that any particular understanding of arbitrariness creates false positives. However, we'll see that it is still possible to show that certain characterizations of arbitrariness create a serious *risk* of false positives.

Conventions' arbitrariness is closely connected to the idea that we conform to conventions *because others conform to them*. Lewis captures this point by arguing that conventions involve coordination by precedent. We have multiple, equally good options—there are alternatives—and we all end up choosing the same option because we have encountered past instances of others choosing that option and we therefore expect others to make the same choice again in the future. As Lewis says, 'the ... source of mutual expectations is precedent'.⁵

Millikan's way of capturing the point is to say that conventional behavior has 'little tendency to emerge or reemerge in the absence of precedent'.⁶ On her picture, past instances of people engaging in the behavior cause new instances of the behavior. Without the past instances of similar behavior, the behavior would not be any likelier to occur than any of its alternatives. Because of their arbitrariness, conventions are deeply dependent on the fact that others have previously engaged in that sort of behavior.

Lewis and Millikan agree that conventions involve perpetuating a practice because others have done it in the past.⁷ As Millikan notes, the main difference between her and Lewis is that Lewis's 'because' is the 'because' of practical reason, whereas Millikan's is purely causal.⁸ In other words, Lewis thinks that precedents justify and motivate us when we conform to conventions, whereas Millikan thinks they merely cause us to conform. But regardless of whose interpretation of the 'because' we accept, arbitrariness is crucial to understanding conventions' dependence on past conformity. It is only because conventions are arbitrary that their persistence is so dependent on the fact that others have previously engaged in that practice. If conventions were practices without alternatives, we wouldn't need the force of conformity to get us to choose the same action—we would just use our only option.

Now let's consider two closely related practices that will help us see how difficult it is to identify the precise sense in which conventions are arbitrary. First, consider nurses' disaster triage practice of checking the pulses of unconscious victims. Second, consider their practice of using the color red to label victims whose treatment should be the highest priority. Both practices are widespread among nurses. The practices spread in a similar manner: during disaster triage training, nurses learn that they should always check for a pulse and that they should always use red to label the patients most in need of care. Despite these similarities, using red to label top-priority patients is clearly conventional, whereas checking for a pulse is not.

What difference between these practices makes one of them conventional and the other not? Checking for a pulse is crucial to determining whether an unconscious person is alive. But using the color red is optional: nurses could just as well use a different color, or a shape instead of a color, to

⁵Lewis, *Convention*, 36.

⁶Millikan, *Language: A Biological Model*, 7.

⁷Others have picked up on this feature of conventions as well. See Marmor, *Social Conventions*, 4 and Miller, *Social Action*, 118. Gilbert says that we follow conventions because they are the way *we* do things, in the much stronger sense of a collective fiat. See Gilbert, 'Rationality, Coordination, and Convention', 17–8.

⁸Millikan, *Language: A Biological Model*, 55–6.

quickly label the most urgent patients. It seems plausible that the difference between these two practices—the feature that makes using red conventional and checking for a pulse not conventional—is arbitrariness. Using red has alternatives; checking for a pulse does not. Because arbitrariness seems to be the most plausible differentiating feature for these two cases, an understanding of arbitrariness that did not differentiate them would create a serious risk of a false positive.⁹

Interestingly, the most straightforward characterization of arbitrariness doesn't actually distinguish the two practices. If we just say that conventions must have alternatives, checking for a pulse will likely count as conventional, too. In the broadest sense nurses have many alternatives to checking for a pulse: they could start clearing away rubble instead of assessing unconscious victims, or leave the disaster scene altogether and go home. In fact, in this very broad sense, there are alternatives to all of our voluntary actions. We could always abandon what we're doing to do something else. That isn't the kind of arbitrariness that is distinctive of conventions.

One step in the right direction is to require that a convention's alternatives be other ways of doing whatever the convention itself is a way of doing. In order to make checking unconscious victims for pulses conventional, an alternative would have to be another way of rapidly determining whether an unconscious person is alive, not just any other possible activity. Similarly, to count as an alternative to wearing a coat, a practice would have to be a way of keeping warm outside in cold weather. A convention's alternatives have to be alternative ways of doing the *same thing*.¹⁰

But of course there actually *are* alternative ways to rapidly determine whether unconscious people are alive: nurses could check for breathing, or shake people to see whether they wake up. However, these alternatives do not render the practice of checking for pulses conventional in the way that the possibility of using another color or a shape makes labeling victims with red conventional. A convention's alternatives are *approximately as good* as the convention.¹¹ Checking for breathing is a significantly worse alternative because some of the most urgent victims might have a pulse but not be breathing. Shaking the victims is still worse. On the other hand, many other colors would work just as well as red for marking the most urgent victims.

So, an alternative is another possible, approximately equally good way of doing whatever the convention does. But what does 'possible' mean in this context? If we read it in the broadest sense, logical possibility, we run into trouble again. There are certainly equally good, logically possible alternatives to checking for pulses, such as instantaneously sensing whether someone's heart is beating by sight. Surely almost everything we do admits of an equally effective, logically possible alternative, so other false positives may crop up as well.

The problem persists if we switch to nomological possibility—that is, to what is possible given how our natural world works. Then we would have to admit that there are equally good alternatives to checking victims for pulses, such as using an electronic device that rapidly scans a

⁹We will work with the triage examples extensively. But there are many other pairs of practices, one conventional and the other not, that share so many features that arbitrariness must be what makes the difference. Consider: planting corn in the spring (non-conventional) versus planting it in rows (conventional) among farmers, deep abdominal breaths (non-conventional) versus moving the diaphragm to create vibrato (conventional) among singers, and cooking chicken before serving it (non-conventional) versus serving chicken in the form of a fried nugget (conventional) among fast food cooks.

¹⁰For more about this notion of a *way of* doing something, see Israel, Perry, and Tutiya, 'Executions, Motivations, and Accomplishments'. There are some thorny issues here that we have not acknowledged: how can we tell exactly what any given practice is a way of doing, and what has to be the case for a practice's alternative to be another way of doing the same thing? These interesting puzzles for action theory exceed the scope of the present discussion.

¹¹Here 'good' is meant in the sense of good at (or for) doing whatever the convention does, and not in a morally loaded sense.

large area for unconscious, living humans. Surely this nomologically possible alternative does not make checking for pulses arbitrary in the way that using red to label top-priority patients is arbitrary.

The solution to this problem may seem obvious: turning to epistemic possibility. It is tempting to think that what matters is not what's actually possible in any sense, but rather what the participants in a convention *know* to be possible.¹² On that understanding of possibility, conventions are practices in which the participants are aware that what they do is arbitrary. The nurses are all aware that there are plenty of ways they could have labeled top-priority victims aside from using the color red, but they do not know of any equally good alternatives to checking for pulses. We can finally distinguish the two cases.

This is Lewis's way of avoiding the risk of false positives. He explicitly requires that participants in a convention know about its alternatives. In fact, he makes the stronger claim that participants in a convention must have *mutual* knowledge of its alternatives. That is, in order for English to count as a conventional language, I must know that you know that I know (and so on, indefinitely) that our linguistic practices have alternatives.¹³

Millikan also leans toward requiring participants' knowledge of alternatives. She writes:

Yet an arbitrariness of pattern in relation to function is not sufficient evidence for conventionality.... That other reproduced patterns would serve just as well if only their possibility was *known* is also not an argument for conventionality. Lighting fires by rubbing two sticks together may be a sort of convention among boy scouts, but it is not merely conventional in a culture that knows no other way to light fires.¹⁴

Here Millikan seems to suggest that lighting fires with sticks is conventional among boy scouts and non-conventional within the culture she imagines because the boy scouts know about other ways of lighting fires, but people in her imagined culture do not. This is plausibly interpreted as implying that knowledge of alternatives is required for conventionality.

However, requiring knowledge of alternatives creates false negatives: it excludes many unrecognized conventions. Tyler Burge makes this point by asking us to imagine a completely isolated linguistic group that is unaware of any alternatives to their own ways of speaking. Their language develops in exactly the same way as other natural languages—the only difference is a single, widespread false belief. It seems bizarre to suggest that this false belief makes their language non-conventional, and therefore fundamentally different from all other languages on the planet, rather than just making them wrong about the nature of their language.¹⁵ We can make the same point using an actual unrecognized convention. Our ignorance of the alternatives to our practices with intonation means we are mistaken; it doesn't make those aspects of our language non-conventional.

Now consider the practice of having mostly men in key decision-making roles (such as CEO) and mostly women serving in auxiliary roles (such as administrative assistant). This is our way of dividing up the jobs that are available in the workplace. If conventions required awareness of alternatives, then when we realized that there are equally good alternatives to this practice (i.e. that women are just as capable as men in decision-making roles), we *made* those practices conventional rather than *realizing* that they were conventional. Bizarrely, the first individuals who became aware of

¹²Really this would be a turn to the intersection of nomological and epistemic possibility: genuine alternatives to a convention must be both nomologically possible *and* believed to be possible. There is an interesting question about whether it might be better to use pure epistemic possibility and count practices that lack actual alternatives but are *believed* to have alternatives as arbitrary in the way characteristic of conventions. I leave that issue aside.

¹³Lewis, *Convention*, 60, 74–6.

¹⁴Millikan, *Language: A Biological Model*, 7.

¹⁵Burge, 'On Knowledge and Convention', 250.

those alternatives should have seen their mission as transforming non-conventional practices into conventional ones. This implies that the practice changes in nature as soon as the participants become aware of a part of its nature, which does not fit the phenomenology of realizing that some part of our lives has actually been conventional all along. As Burge puts it, to say that a practice becomes conventional only once we recognize its alternatives ‘would be to use ‘conventional’ unconventionally’.¹⁶

These intuitive considerations are not decisive, but they do increase the desirability of eliminating the risk of false positives without appealing to participants’ awareness of alternatives. There is more to Lewis’s and Millikan’s attempts to walk the tightrope between false positives and false negatives than I have acknowledged so far, and we will return to them in Section IV. But first I’ll offer a positive suggestion: a way of eliminating the risk of false positives and ruling in all of the unrecognized conventions that avoids appealing to epistemic possibility.

III. The Balance Pole: Accessibility

We’ve seen that simply saying ‘conventions must have alternatives’ creates a serious risk of false positives. We’ve also seen that saying ‘participants must know about the alternatives’ creates false negatives (specifically, many unrecognized conventions are wrongly excluded). I’ll argue that the solution to this problem is the Accessibility Requirement developed below. First, we’ll work toward the right notion of accessibility in Section III.i. We won’t spell out the Accessibility Requirement for conventions until Section III.ii, but nonetheless we will construct our notion of accessibility with reference to its ability to help us understand conventions. Accessibility is a phenomenon that goes beyond conventions, but the particular kind of accessibility we need is constrained by the role it plays in conventions.

III.i. What is accessibility?

Let’s think again about nurses’ practice of checking for pulses. Why can we be sure that this practice is non-conventional, despite its nomologically possible, equally good alternatives? Because those alternatives are not ones the nurses can *really* use. The nomological possibility of inventing a device that rapidly scans a large area for unconscious, living humans doesn’t do us any good because, we might say, that alternative is not very *accessible* to us. On the other hand, why is Burge’s isolated group’s language conventional, despite their ignorance of its alternatives? Because their ignorance doesn’t change the fact that the alternatives could easily have been used instead: very small differences in the circumstances of the group at the time when their language was developing could have led to an entirely different set of communicative practices. Those alternatives were real, if forgotten, threats to the dominance of their linguistic practices because they were so easily *accessible* to the group at the time when their language began to emerge.

Burge’s positive suggestion for how to understand the arbitrariness of conventions is along these lines. He writes, ‘As a matter of fact—whatever the participants may believe—it is within the power of the participants to have learned an incompatible regularity that would have served substantially the same social functions without demanding significantly greater effort on the part of the participants’.¹⁷ Burge offers this suggestion at the end of his short critical piece about Lewis, so

¹⁶Ibid.

¹⁷Ibid., 254. Marmor also briefly discusses a similar sort of constraint on a convention’s alternatives. See Marmor, *Social Conventions*, 9–10.

he does not provide any development of this idea. In fact, we will need to go well beyond Burge's remark to arrive at a full understanding of the arbitrariness of conventions. Nonetheless, as we move forward, Burge's idea of a convention's alternatives not requiring *significantly greater effort* than the convention itself will be important.

Notice how natural it was to describe the equally good alternatives to checking for pulses as 'not very accessible' and the alternatives to Burge's imagined language as 'easily accessible'. This suggests that accessibility is a matter of degree: one alternative can be more or less accessible than another. For example, when I poach an egg on my stove, the alternative of poaching it in a microwave oven is at least as accessible. But the alternative of poaching the egg in a pot over an open fire is much less accessible. I would have to gather wood and build a fire, which takes significantly more effort than using my kitchen. There is a spectrum of accessibility, ranging from the extremely accessible to the completely inaccessible.

The egg example also highlights the fact that claims about an alternative's accessibility must specify to *whom* the alternative is accessible. Poaching an egg on a stove and in a microwave are very accessible to me because I own those appliances, but they are much less accessible to someone who does not. For such a person, poaching eggs over an open fire may be the most accessible option.

In light of these insights, here is my definition of 'accessibility':

Accessibility

An option's level of accessibility to an individual or group is the degree of ease with which that individual or group can encounter and use it.

'Option' here refers to a way of doing something that can be done in multiple ways. Options then count as alternatives relative to each other, provided that they are ways of doing the same thing (e.g. flying from Los Angeles to San Francisco is an alternative to making the journey by car, and vice versa).¹⁸ Accessibility does not require any sort of knowledge or awareness. An option's level of accessibility to someone is just a matter of how feasibly she could happen upon and employ it, whether she is aware of this or not. We'll focus attention on the definition's key elements in turn.

When we discuss the ease with which a group or individual *can* encounter and use an option, 'can' should be understood in the sense of nomological possibility. An option's accessibility is a matter of how easily the relevant group or individual can encounter and use that option, within the basic constraints of our natural world. Is it nomologically possible for them to encounter and use that option with very little effort? Then it is very accessible. Otherwise, it is less accessible. Accessibility, then, involves a restricted form of nomological possibility.

The notion of *encountering* an option requires clarification. One way to encounter an option is to be exposed to an instance of that option. For example, if I travel to a place with different gender practices, I encounter those alternatives even if I don't realize it. I can also encounter an option by hearing about it from someone else, or by thinking of it on my own. To *use* an option is just to behave in a way that is an instance of that option: to use the option of driving from Los Angeles to San Francisco is to complete that drive on a particular occasion.

Our definition of 'accessibility' applies to groups as well as individuals, but the preceding discussion of encountering and using options applies only to individuals. Given that groups are made of individuals, a group presumably encounters or uses an option in virtue of its members encountering or using that option. It seems too demanding to require that *every* individual in a group encounter or use an option in order for the group to have encountered or used it, but it seems too

¹⁸Really, accessibility is a broader notion that applies to any practice, not just to practices that count as options in my terminology. The definition has been restricted to practices that count as options for the sake of simplicity.

permissive to say that the group as a whole has encountered or used an option just because a single member encountered or used it.

I propose a ‘majority rules’ approach: a group has encountered (or used) an option just in case most members have encountered (or used) it. If some group members, but not the majority, have encountered an option, we can say that the option has been encountered within the group, but not that the group itself has encountered the option. When Henry Heimlich invented the Heimlich maneuver, he made it the case that the option had been encountered *within* the group of United States citizens, but not until most people in the country had encountered it could we say that US citizens as a group had encountered it.¹⁹

Similarly, a group *uses* an option just in case the majority of the group’s members sometimes use it. We needn’t require that all of them use it, or that any of them use it all the time. It’s true that US citizens, as a group, use the option of watching television, even though not all of them watch television, and none of them do it all the time. This approach allows us to capture the right type of threat to a convention: an option that a few members of a group could easily encounter and use is not very threatening, but an option that *most* members could easily encounter and use surely is.

Now we’ll zoom in on *ease*. This is the part of accessibility that admits of degrees. It is also where Burge’s insight about effort comes in. For an individual, the ease of encountering and using an option is a matter of how much effort she would have to expend to encounter and use it. For a group, it’s determined by the total amount of effort that all group members would have to expend in order for most of them to encounter and use the option. Accessibility correlates inversely with required effort. At one end of the accessibility spectrum are completely inaccessible options: options the group or individual could never encounter and use. At the other end of the spectrum are options that require very little effort, such as wiggling our fingers. Between these two extremes lies a wide range of required effort for encountering and using various options, which gives rise to a wide range of accessibility.

Now that we’ve discussed our key terms, we need to examine the idea that an option’s accessibility is the degree of ease of encountering *and* using it. An individual or group’s level of ease of encountering an option can diverge from their level of ease of using the option once they encounter it, which is why it is crucial to add this distinction to Burge’s initial insight about effort. The level of ease of encountering *and* using an option is equivalent to the level of ease of encountering it or the level of ease of using it, whichever is lower. The level of accessibility is determined by the ease with which the group or individual can do both, so it is constrained by the lower level of ease. To illustrate this point, we’ll consider some cases.

When ancient Egyptians first began using sundials, there were people living without sundials in, for instance, Northern Europe. If Northern Europeans had somehow encountered sundials, using that option would not have taken a huge amount of effort. Sundials are not terribly difficult to construct or read. However, encountering a sundial was no easy task for them. The journey to Egypt would have been arduous; inventing sundials would have required a significant technological advance. Overall, this option does not seem to have been very accessible to those Northern Europeans. The fact that they could easily have used that option if they encountered it did them no good, and presented no threat to their established ways of telling what part of the day it was, because it was so difficult for them to encounter it. The lower value—their ease of encountering the option—determines the option’s level of accessibility to them.

Now imagine that, in the future, one of the smartest humans develops an amazing new method for speed-reading that requires rapidly integrating and extrapolating from large amounts of

¹⁹It’s worth noting that the Heimlich maneuver was once an unrecognized convention. Many people thought it was the best way to clear a choking person’s airway, but actually sharp blows to the back are an equally good alternative.

information. Her demonstrations of the method go viral on the Internet. It would be pretty easy for Internet users to encounter the speed-reading option once it went viral. However, the method is extremely difficult for even humans of above-average intelligence to use. As a result, it would not be very accessible to Internet users. Although most of them could pretty easily encounter the alternative, it's still not something that could realistically compete with their usual ways of reading because it is so hard to use. An option's level of accessibility is determined by the ease of encountering it or the ease of using it, whichever is lower.

III.ii. The Accessibility Requirement

Now that we have a solid grasp on accessibility, we are ready to develop the Accessibility Requirement. To see what shape it should take, we'll return to the two ways of falling off the tightrope: risking false positives and creating false negatives.

First, the potential false positives. We can't eliminate the risk just by requiring that a convention's alternatives be accessible. Consider again the triage practice of checking for pulses. Some of the equally good alternatives to checking for pulses are presumably somewhat accessible to nurses. Medical researchers could probably encounter other equally good ways of determining whether an unconscious person is alive with a large but finite amount of effort. Nonetheless, checking for pulses is non-conventional for the following reason: its equally good alternatives are significantly *less* accessible. They are not genuine threats to the current practice of checking for pulses because they are significantly more difficult for nurses to encounter and use. So, we should require that a convention's alternatives be at least as accessible as the conventional practice itself rather than trying to specify an absolute minimum threshold of accessibility.

Next, let's return to the false negative cases, our unrecognized conventions. So far, we are still in danger of excluding them. Consider gender conventions. Once these practices are deeply entrenched, their alternatives probably *are* less accessible than the conventional practices. Within many groups during the nineteenth century, the convention of wives staying home to maintain the household while husbands worked outside the home was very entrenched, and as a result it was more accessible than its alternatives. To work outside the home, a woman would have had to overcome the disapproval of her family, convince someone to hire her, and possibly struggle to obtain the necessary education. But importantly, before this convention got going, there were plenty of equally accessible alternate ways to divide up labor within families: wives could have gotten jobs instead of husbands, or each spouse could have spent time working both outside and inside the home.

Similarly, consider Burge's isolated linguistic group. At the time when their communicative practices first began, they could just as easily have developed different linguistic conventions. But once their communicative practices became well established, those practices were more accessible to them than the alternatives. Crucially, those options became more accessible only *because* they were the group's established practices—they didn't become immune to the threat posed by their alternatives until they had already overcome that threat in the paradigmatic conventional manner. So, it seems we should index the equal accessibility of a convention's alternatives to the time immediately *before* one of the options took off and became conventional.

It seems likely that this is what Burge was getting at when he wrote, 'it is within the power of the participants to have learned an incompatible regularity that would have served substantially the same social functions without demanding significantly greater effort'. Even though *now* using a different language would require significantly more effort, what matters is that back at the time right

before their current practice took off, the participants *could have* just as easily adopted a different, equally good practice.

However, saying that a convention's alternatives must have been equally accessible immediately before the convention took off would create another set of false negatives. Consider the practice of sending invitations to formal events by means of the postal service. This was once the most accessible way to invite people to a formal event. But then email became equally, and eventually more, accessible to many groups of people. Now it is *conventional* to invite people to formal events, such as weddings or graduation ceremonies, by sending a physical letter. If we require the existence of equally good and accessible alternatives *before* the practice is established, we will exclude this convention. At the time when the practice of sending invitations by mail began, there were no equally accessible, equally good alternatives.

More generally, saying that a convention's alternatives must have been equally accessible before the convention began would exclude practices that shift from non-conventional to conventional without an interruption in the practice's prevalence. This shift can happen when a new option becomes equally accessible, as email did in our example above. Knitting by hand is another example of such a practice. It was once the best way to obtain a warm sweater, but obtaining a sweater made by industrial machinery is now more accessible to many groups than knitting a sweater by hand. Nonetheless, knitting by hand has continued to be prevalent in some of those groups, where it has become a conventional practice. We are in danger of excluding this convention as well.²⁰

To solve this problem, we will have to avoid indexing the accessibility of a conventional practice's alternatives to the time before the practice began. Instead, we can say that a convention's alternatives must be as accessible as the conventional practice itself, independent of the dominance that the conventional practice gained as it became conventional—or, more precisely, the dominance it gained as people did it because others were doing it. If the only reason a widespread practice is more accessible than its equally good alternatives is that people are conforming to what others are doing, the practice can still be conventional. But if a practice is more accessible than its alternatives for other reasons, such as in our checking pulses example, it can't be conventional. Conventions are practices in which we engage because others have done so, and doing something in the conventional way *is* often easier—but only because that's how so many other people have done it.

This understanding of the accessibility of a convention's alternatives accommodates our new false negative cases (i.e. the formal invitations and knitting by hand) as well as our original false negative cases (i.e. Burge's language case and the gender conventions). Our practice of using the postal service to send formal invitations *does* have an equally good, equally accessible alternative: email. This is true even before we factor out the dominance that using the postal service has gained due to the fact that others have done it. The same goes for knitting by hand, so neither of these practices is excluded. And although the alternatives to the linguistic practices of Burge's isolated group eventually become less accessible to them than the practices they use, this is only because of the dominance their practices gained as they became conventional. Independent of that dominance, the alternatives are equally accessible, so that practice is not excluded from conventionality either. We can say the same for gender conventions such as the convention of wives staying home while husbands work outside the home.

We are finally ready for a precise formulation of the requirement that conventions must have alternatives, which includes the Accessibility Requirement.

²⁰Jody Azzouni also discusses cases in which a practice shifts from non-conventional to conventional. See Azzouni, 'Defense of Logical Conventionalism', 36.

Alternatives Condition for Conventions

For doing *B* as a way of doing *A* to be conventional in a group, there must be at least one other way of doing *A* that is at least approximately as good and at least approximately as accessible to the group as doing *B*, independent of the dominance doing *B* gained as people did *B* because others did *B*.

The Alternatives Condition expresses the sense in which conventions are arbitrary. It is a necessary, but not sufficient, condition for conventionality. The Accessibility Requirement is the part that requires that a convention's alternatives be 'at least approximately as accessible to the group as doing *B*, independent of the dominance doing *B* gained as people did *B* because others did *B*'. Two other previously discussed requirements are built into the Alternatives Condition: the requirement that a convention's alternatives be other ways of doing whatever the convention is a way of doing, and the requirement that the alternatives be at least approximately as good as the convention.

We should say more about the idea that a convention's alternatives must be 'at least approximately as accessible' as the convention. The 'at least' part of this requirement allows a convention's alternatives to be far *more* accessible than the conventional practice itself. For example, various groups of people have followed complex eating conventions, such as using separate forks for meat, fish, salad, etc. Some of the alternatives to those practices are actually far more accessible than the practices themselves. The option of eating an entire meal with just one fork is far easier to encounter and use than the option of using multiple forks for highly specific purposes.

The 'approximately' part of the requirement is a little trickier because it allows a convention's alternatives to be slightly less accessible than the convention. This is because alternatives that are slightly less accessible than the current practice might nonetheless have been genuine threats to the practice's dominance: they could *quite easily* have been used instead. One might object to the vagueness this brings into the Accessibility Requirement. However, this vagueness is benign because conventions exhibit similar vagueness. Whether a practice is arbitrary enough, or widespread enough, to count as a convention is sometimes indeterminate. What's important is that practices whose only alternatives are significantly less accessible fail to depend on past conformity in the right way. As Millikan would put it, if those practices disappeared, they would be significantly likelier to reemerge than their alternatives because they are significantly easier to discover and use.

We should also note that the 'because' in the 'independent of the dominance doing *B* gained as people did *B* because others did *B*' clause is neutral between Lewis's rational construal and Millikan's causal one. In a complete account of conventions, it would be filled in with something more specific about the type of influence past uses of *B* exert on present ones, but a complete account of conventions is not our goal.

The Alternatives Condition properly categorizes our original triage cases, along with the false negatives we have already discussed. Checking for pulses during disaster triage (a potential false positive case) does not satisfy the Alternatives Condition. All of the alternatives that are at least as accessible as checking for pulses (such as checking for breathing) are inferior. And all of the alternatives that are at least as good as checking for pulses are far less accessible. Even if we factor out the dominance checking for pulses acquired due to the fact that others had done it, the equally good options would have required significant scientific advances beyond what was required for just feeling for a pulse. On the other hand, using red to label the most urgent patients *does* satisfy the Alternatives Condition: independent of the dominance red has now gained, many other colors and shapes are just as good and just as accessible.

In fact, given how neatly the Accessibility Requirement categorizes our cases, one might wonder whether it is too ad hoc, and whether it fits those cases at the expense of miscategorizing others. To stave off that concern, let's consider some paradigmatic conventions. First, one of

Lewis's classic examples: the convention of driving on the right side of the road in, for instance, the United States.²¹ This practice has one equally good alternative: driving on the left. Notice that this alternative is equally good only independent of the dominance driving on the right gained as it became conventional, because of course if we take that dominance into account, driving on the right in the US is far superior to driving on the left from a safety standpoint.²² Furthermore, driving on the left is equally accessible (again, independent of the acquired dominance of driving on the right). So, the practice of driving on the right satisfies the Alternatives Condition.

Now let's consider a paradigmatic unrecognized convention. Despite the fact that few people realize it, the standing distance that we maintain during conversation is conventional.²³ Every society's standing distance practice has equally good alternatives: all of the standing distances that would allow conversational partners to talk without having to yell to be heard and without being so close that their eyes cross. These alternatives are, independent of the practice's acquired dominance, just as accessible as the standing distance the society uses. No standing distance is harder to discover than the others. So, these practices satisfy the Alternatives Condition too.

IV. Why Lewis and Millikan Can't Walk the Tightrope

Now that we've found a way to walk the tightrope between false negatives and the risk of false positives, I'd like to do more to convince you that neither Lewis nor Millikan can make it across. We saw in Section II that both authors seem tempted to require knowledge of alternatives, thereby creating false negatives. In both cases there is more to the story. Lewis attempts to accommodate unrecognized conventions by relying on tacit or merely potential knowledge of alternatives. And there is some evidence that Millikan may actually have been thinking along the lines of the Accessibility Requirement herself. Below I'll show that neither of their approaches works, at least not without major changes to their accounts. The Accessibility Requirement goes beyond what Lewis and Millikan can do.

First, we'll consider Lewis. We'll reproduce his account and then focus on the part that threatens to exclude unrecognized conventions. Lewis writes:

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*....
- (5) *R* is not the only possible regularity meeting the last two conditions....
- (6) Finally, the various facts listed in conditions (1) to (5) are matters of *common* (or *mutual*) knowledge.²⁴

²¹Lewis, *Convention*, 40–1.

²²This example shows why it is important that not just the equivalent accessibility but also the equivalent goodness of a convention's alternatives fall within the scope of the 'independent of the dominance doing *B* gained as people did *B* because others did *B*' clause. Millikan makes a similar point. See Millikan, *Language: A Biological Model*, 57.

²³Millikan, 'Deflating Socially Constructed Objects', 34.

²⁴Lewis, 'Languages and Language', 5–6.

To summarize: a convention is a regularity to which there is an alternative and to which nearly everyone in a population conforms because she believes that others conform, and they all prefer that everyone else conforms, and all of this is mutually known among them.

We'll focus on conditions (5) and (6). Condition (5) builds in alternatives, and condition (6) requires not only that participants know that all of the preceding conditions obtain, but also that this fact be mutual knowledge among them. In order for driving on the right side of the street to be conventional, I must know (among other things) that there is an alternative, that you know there is an alternative, that you know that I know there is an alternative, etc.²⁵

Lewis provides two options for accommodating unrecognized conventions. First, he suggests that our knowledge of conventions can be 'merely potential' in the sense that we all 'have evidence from which we could reach the conclusion' that there is a convention, but perhaps we never do. Second, he suggests that this knowledge may be actual, but 'irremediably nonverbal' or 'tacit'.²⁶

These strategies are promising for accommodating gender conventions. Before they realized that many gender practices are conventional, most people had likely met others who did not completely conform. Many had probably also experienced discomfort with their own gender roles. Lewis could say that most people either had evidence that gave them potential knowledge of the alternatives or already had that knowledge in an unconscious and repressed way.

However, this solution will not work when participants lack evidence of their convention's alternatives, such as in Burge's isolated linguistic group. Because they are isolated, the group will not encounter alternatives to their own language, so they will lack evidence from which to conclude that their communicative practices have alternatives. And because they lack this evidence, we lack a basis for claiming that they have potential or tacit knowledge of their practices' alternatives.

Generally, Lewis runs into trouble with conventions that have been around so long that the alternatives that originally were a threat have been entirely lost. For a second example, imagine a group that long ago began the practice of eating no fish but trout, despite the fact that other kinds of fish were plentiful and equally nutritious. Perhaps they initially wanted to build group identity with distinctive dietary practices. Decades later, the group has become isolated. Everyone now believes that other kinds of fish are poisonous and never tries to eat them. Eating trout is an unrecognized convention in that group, despite their lack of evidence of the equally good alternatives, and this is something Lewis is unable to capture. Lewis's account still creates false negatives.

Next, we'll turn to Millikan's account of conventions. Millikan argues 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.²⁷

Condition (1) is where Millikan specifies that a convention persists *because* (in the causal sense) others did it in the past. Condition (2) is where she incorporates equally good alternatives by saying that if a convention has a function (i.e. if it serves some further purpose), there must be other equally good options for performing that function.

²⁵This is a simplification of Lewis's notion of mutual knowledge, but for our purposes it is not misleading. See Lewis, *Convention*, 54.

²⁶Lewis, *Convention*, 63–4; cf. Lewis, 'Languages and Language', 25.

²⁷Millikan, 'Deflating Socially Constructed Objects', 31.

In this formulation of Millikan’s account, which is more recent than the one from which we quoted in Section II, there is no risk of excluding unrecognized conventions. It might seem that all that’s needed is to add the Accessibility Requirement to rule out the risk of false positives. In fact, Millikan hints at a notion similar to accessibility, so it might even seem that the Accessibility Requirement is implicitly part of her view. Others have made suggestions to that effect.²⁸ Millikan’s account is clearly a response to the complexity of Lewis’s analysis and the problems it generated, including the problem Burge identified.²⁹ She may well have had in mind something like Burge’s way of spelling out the arbitrariness of conventions, even though she didn’t explicitly build it into her account.

However, the Accessibility Requirement is incompatible with Millikan’s account. The incompatibility stems from Millikan’s avoidance of indexing conventions to groups. She sees conventionality as just a feature that a type of behavior gains due to how it is reproduced. Sometimes each *instance* of reproduction requires a group (consider the convention of dancing around a maypole), but that does not imply that there is one overarching group in which the convention obtains. Millikan does not say exactly why she eschews groups, but we can infer her reason from one of her objections to Lewis’s account.

As we saw above, Lewis thinks that conventions exist within groups (or, in his terminology, populations). Millikan objects that it is often impossible to independently identify the group in which a convention exists. She discusses the convention of using ‘Break a leg’ to mean the same as ‘Good luck’. There is no immediately obvious way to identify the group that uses the saying. If we stipulated that the relevant group is simply the people who actually do use the saying, that would ensure that we’ve identified a group that uses it, but it would also suggest that arbitrary behavior can be construed as conventional among whoever practices it.³⁰ So, an account of conventions that indexes them to groups is too permissive. Millikan avoids this problem: she says that conventions are just types of behavior passed on in a certain way. Groups don’t matter.

On the hypothesis that Millikan had something like the Accessibility Requirement in mind, let’s try explicitly adding it to her account: a pattern of activity is conventional if it is reproduced and it has alternatives that are at least as good at serving the convention’s function and at least approximately as accessible, independent of the dominance the convention gained as people reproduced it. Now, this account is incomplete because accessibility is always accessibility *to someone*. Since Millikan doesn’t want to bring in groups, she could say that the alternatives must be equally accessible to the particular participant(s) in each reproduced instance of the conventional pattern.

The problem with this approach is that the chain of reproducing a pattern of behavior in the conventional manner could then get broken quite easily. Imagine a town in which it is conventional to ride a bicycle to a creek several miles away to go swimming. There are paved roads that go to the creek, but hardly anyone ever drives there; riding bikes to the creek is a time-honored tradition. Now consider a 10-year-old resident of the town who rides a bicycle to the creek with her friends. She has

²⁸Specifically, Millikan says that a conventional pattern of activity is ‘no more able than other known or *equally knowable* patterns to serve its function’. Millikan, *Language: A Biological Model*, 32 (my italics). Brian Epstein attributes to Millikan the insight that ‘conventionality involves the “availability of genuine alternatives”’ but notes that ‘this itself isn’t a straightforward notion to cash out’. Epstein, ‘Review of Ruth Millikan’, 5. Similarly, while spelling out Millikan’s account of conventions, Jody Azzouni writes: ‘technological development can induce conventionality because it creates *practical* alternatives that weren’t there before’, but ‘[t]here is a lot of work to do here ... detailing more fully the notion of “genuine practical alternatives”’. Azzouni, ‘Defense of Logical Conventionalism’, 36. The present paper does the needed work of identifying accessibility as what makes an alternative genuine, but also shows (below) that it moves us away from a central aspect of Millikan’s view.

²⁹Millikan, *Language: A Biological Model*, 1–2.

³⁰*Ibid.*, 29.

no equally accessible, equally good alternatives; she doesn't know how to drive a car, and walking would take too long.

Would the child's behavior constitute a break in the chain of conventionality according to our revised version of Millikan's account, because the alternatives to the child's behavior are not accessible *to her*? And if the only link between past instances of a behavior and present ones is the behavior of an individual to whom the alternatives were inaccessible, does the old convention end and a new one begin? Or could we count the new behavior as an indirect reproduction of the old conventional behavior? In the latter case, Millikan would be able to say that the convention persists, but she would have to admit that children cannot fully participate in some conventions. Millikan would be reluctant to accept this outcome: she is explicitly concerned with allowing for full participation in conventions by children, contra Lewis.³¹

Untying this knot is easy if we index conventions to groups. Then we can say that the *group* is the entity to which a convention's alternatives must be accessible, and the child's behavior of riding her bicycle to the creek counts as conventional because she is a member of a group in which there are genuinely threatening alternatives to that way of doing things. Only by indexing the accessibility of alternatives to groups can we capture the participation in conventions of children and others to whom the alternatives are not individually accessible (such as, perhaps, adults with disabilities). And aside from the fact that Millikan is explicitly committed to capturing the participation of such individuals in conventions, it is important in its own right: children, and adults with disabilities, certainly appear to participate in the conventions practiced within their broader social groups. Excluding them would require some serious justification.

Thus, the Accessibility Requirement is not compatible with Millikan's account of conventions: it works only if we index conventions to groups, and Millikan is unwilling to do so. Millikan could either index conventions to groups in order to make use of the Accessibility Requirement,³² or she could search for a way to index accessibility to individuals without excluding children (and adults with disabilities) from participation in conventions. The former would be a move away from a central feature of her account, and the latter promises to be a difficult task.³³

Neither Lewis nor Millikan can walk the tightrope between false negatives and the risk of false positives. However, with the Accessibility Requirement in hand, we can. This means we are well positioned to work toward a complete account of conventions that incorporates the precise sense in which they are arbitrary, without excluding the subtle, unrecognized conventions that shape our lives.

³¹Ibid., 55–6.

³²Of course, this option would not be viable if Millikan's objection to indexing conventions to groups were sound. But, returning to Millikan's example, all we would have to do to identify the group in which saying 'Break a leg' is conventional is to identify all the people who are involved with theater or interested in it. The boundaries of such groups are often quite vague, but again, that vagueness is part of the nature of conventions rather than a symptom of error.

³³Perhaps the pressure to index conventions to groups is the reason why Millikan doesn't go beyond her hint in the direction of the Accessibility Requirement. When she remarks that a convention's alternatives must be 'equally knowable', she never asks: equally knowable to whom? Millikan, *Language: A Biological Model*, 32.

Acknowledgements

For helpful conversations related to this paper, I am grateful to Daniel Ehrlich, Peter Graham, Meredith McFadden, Ruth Millikan, Luis Montes, Courtney Morris, Michael Nelson, John Perry, John Ramsey, Patrick Ryan, Howard Wettstein, and Monique Wonderly, as well as audiences at the 2013 Pacific APA, the University of Western Ontario's 2013 Graduate Conference, and the 2013 Northwest Philosophy Conference. I am also grateful for helpful comments from several anonymous referees.

References

- Azzouni, Jody. 'A Defense of Logical Conventionalism'. In *The Metaphysics of Logic*, ed. Penelope Rush, 32–48. Cambridge: Cambridge University Press, 2014.
- Bennett, Jonathan. *Linguistic Behaviour*. Cambridge: Cambridge University Press, 1976.
- Burge, Tyler. 'On Knowledge and Convention'. *The Philosophical Review* 84, no. 2 (1975): 249–55.
- Cain, M.J. 'Conventions and Their Role in Language'. *Philosophia* 41, no. 1 (2013): 137–58.
- Chomsky, Noam. 'Reply to Millikan'. In *Chomsky and His Critics*, ed. Louise M. Antony and Norbert Hornstein, 308–15. Malden: Blackwell, 2003.
- Epstein, Brian. 'Review of Ruth Millikan, *Language: A Biological Model*'. *Notre Dame Philosophical Reviews* 5 (2006): 1–10.
- Gilbert, Margaret. *On Social Facts*. London: Routledge, 1989.
- Gilbert, Margaret. 'Rationality, Coordination, and Convention'. *Synthese* 84, no. 1 (1990): 1–21.
- Guala, Francesco. 'The Normativity of Lewis Conventions'. *Synthese* 190, no. 15 (2013): 3107–22.
- Israel, David, John Perry, and Syn Tutiya. 'Executions, Motivations, and Accomplishments'. *The Philosophical Review* 102, no. 4 (1993): 515–40.
- Lewis, David. 'Languages and Language'. In *Language, Mind and Knowledge: Minnesota Studies in the Philosophy of Science*, ed. Keith Gunderson, 3–35. Minneapolis: University of Minnesota Press, 1975.
- Lewis, David. *Convention: A Philosophical Study*. Oxford: Blackwell Publishers Ltd., 2002.
- Marmor, Andrei. 'On Convention'. *Synthese* 107, no. 3 (1996): 349–71.
- Marmor, Andrei. *Social Conventions: From Language to Law*. Princeton: Princeton University Press, 2009.
- Miller, Seumas. 'Rationalising Conventions'. *Synthese* 84, no. 1 (1990): 23–41.
- Miller, Seumas. *Social Action: A Teleological Account*. Cambridge: Cambridge University Press, 2001.
- Millikan, Ruth Garrett. *Language: A Biological Model*. Oxford: Oxford University Press, 2005.
- Millikan, Ruth Garrett. 'Deflating Socially Constructed Objects: What Thoughts Do to the World'. In *Perspectives on Social Ontology and Social Cognition*, ed. Mattia Gallotti and John Michael, 27–39. Dordrecht: Springer, 2014.
- Moore, Richard. 'Imitation and Conventional Communication'. *Biology & Philosophy* 28, no. 3 (2013): 481–500.
- Schiffer, Stephen. *Meaning*. London: Oxford University Press, 1972.
- Sillari, Giacomo. 'Rule-Following as Coordination: A Game-Theoretic Approach'. *Synthese* 190, no. 5 (2013): 871–90.
- Tomasello, Michael. 'Conventions Are Shared'. *Symposium on Language: A Biological Model by Ruth Millikan, SWIF Philosophy of Mind Review* 5, no. 2 (2006): 29–36.
- Tummolini, Luca, Giulia Andrighetto, Cristiano Castelfranchi, and Rosaria Conte. 'A Convention or (Tacit) Agreement Betwixt Us: On Reliance and Its Normative Consequences'. *Synthese* 190, no. 4 (2013): 585–618.
- Witek, Maciej. 'An Interactional Account of Illocutionary Practice'. *Language Sciences* 47 (2015): 43–55.