

PDF hosted at the Radboud Repository of the Radboud University Nijmegen

The following full text is a publisher's version.

For additional information about this publication click this link.

<http://hdl.handle.net/2066/191041>

Please be advised that this information was generated on 2019-11-18 and may be subject to change.

ORIGINAL ARTICLE

Convention and common ground

Bart Geurts

Department of Philosophy, University of Nijmegen, Nijmegen, The Netherlands

Correspondence

Department of Philosophy, University of Nijmegen, Nijmegen, The Netherlands.
Email: brtgrts@gmail.com

Conventions are regularities in social behaviour of the past that enable us to coordinate our actions. Some conventions are lawlike: they are expected to be observed always or nearly always. However, in order to coordinate our actions, it may suffice that a precedent has occurred often enough, and sometimes even a single precedent will do. So, in general, conventions merely enable us to solve our coordination problems; lawlike conventions are a special case. Grammatical conventions are often lawlike; sense conventions are typically enabling. In order to resolve the indeterminacies that sense conventions give rise to, interlocutors must rely on the common ground. In this and other ways, common ground is a prerequisite for convention-based communication.

KEYWORDS

common ground, communication, convention, speech acts

1 | INTRODUCTION

Syntactic rules, words and their meanings, phoneme inventories: they all vary from language to language, and are therefore a matter of convention. Plato's *Cratylus* dialogue shows that this claim used to be controversial, but that was a while ago; since then there has long been a comfortable consensus in favour of the conventionalist view. This is not to deny that there are non-conventional constraints on the emergence and development of language, but it is to say that they fall far short of determining these processes.

What does it mean for a language to be conventional? Conventions are patterns in social interaction, like driving on the right or shaking hands. The languages studied by phonologists, lexicologists, and syntacticians are abstract, many-layered systems, and it is by no means obvious that such systems qualify as patterns in social interaction on a par with shaking hands or driving on the right. Lewis frames the issue as follows:

This is an open access article under the terms of the Creative Commons Attribution-NonCommercial-NoDerivs License, which permits use and distribution in any medium, provided the original work is properly cited, the use is non-commercial and no modifications or adaptations are made.

Suppose that with practice we could adopt any language in some wide range. It matters comparatively little to anyone (in the long run) what language he adopts, as long as he and those around him adopt the same language and can communicate easily. Each must choose what language to adopt according to his expectations about his neighbors' language: English among English speakers, Welsh among Welsh speakers, Esperanto among Esperanto speakers, and so on. (Lewis, 1969, pp. 7–8)

More formally, let “languages” be mappings from sentences to sentence meanings (e.g., propositions), and let “language” be a social activity. Then Lewis's question is: How are languages and language related? Lewis (1975) seeks to answer this question by defining what it means for a population *P* to “use a language *L*”:

A population *P* uses a language *L* iff there is in *P* a convention of truthfulness and trust in *L*.

That is to say, if *P*-speakers generally avoid uttering false *L*-sentences and *P*-hearers generally accept utterances of *L*-sentences as true, then *P* uses *L* in Lewis's designated sense of the word. These proclivities count as conventional because, in principle, they could have been directed at a language other than *L*.

While Lewis's account has been criticised by Hawthorne (1990) and others for failing to provide a satisfactory answer to the question of when a population uses a language, my worry is with the question itself. For one thing, if we adopt a generative stance on languages, as Lewis does, then it is by no means obvious that speakers use languages at all; for on this view a language is infinitely large, and only a fraction of its sentences are short and simple enough ever to be produced and understood. Apart from that, there are obvious concerns about context dependence, rampant ambiguity, and variation within and between speakers; all of which militate against the idea that people use languages in Lewis's sense. But more importantly, we would like to know what it means for a word, syntactic construction, or stress pattern to be conventionally used in such-and-such a way. Lewis does not tell us that, and does not mean to either; indeed, he denies that the rules of syntax and semantics are conventional (1975, p. 24). Lewis's objective is merely to specify the conditions under which a population may be said to use a language, and his proposal effectively treats sentences as atoms, leaving sub-sentential units out of account.

A more promising approach, I submit, is to proceed bottom up, starting with the phonemes, words, constructions, intonation contours, and so on. Like traffic lights and handshakes, these linguistic building blocks are devices that we conventionally use to regulate our social interactions. They are more like handshakes than traffic lights in that they are instantiated by short-lived actions which may serve as precedents for future actions. Conventions are just precedent types which have gained momentum through frequency of use. Hence, in order to understand how linguistic conventions work, we must first learn to understand how precedents work. But before turning to that topic, I will comply with the standard protocol, in discussions of conventions, and register my main points of agreement and disagreement with Lewis's (1969) celebrated analysis.

2 | CONVENTIONS AND PRECEDENTS

On Lewis's account, the purpose of a convention is to solve a coordination problem, like sharing roads, for example. There are two obvious ways of regulating traffic so as to maximise efficiency and minimise loss of life and property: people driving in the same direction should either keep all to

the left or keep all to the right. It does not matter which solution is adopted, as long as it is the same for all. It is not good enough that most drivers keep to the right; general conformity is best. Finally, everyone's expectation that everybody else drives on the right gives me a compelling reason to drive on the right myself. Generalising from these observations, Lewis proposes the following analysis:

A behavioural regularity R is a convention in a population P iff it is common ground among P 's members that:

- Everyone conforms to R .
- The belief that the others conform to R gives everyone a decisive reason to conform to R himself.
- There is a general preference for general conformity to R .
- R is not the only possible regularity meeting the last two conditions. (Conventions are arbitrary.)

Although practically every part of Lewis's theory has been under attack (Rescorla, 2014), my general impression is that it has withstood the test of time remarkably well, and that its main flaw is that it is too narrow in some respects. Put otherwise, it seems to me that, at the very least, Lewis conventions are an important special case. I will discuss three points on which Lewis's account has been criticised.

First, Lewis's theory of conventions has been lambasted for relying on the notion of common ground, or "common knowledge," as Lewis calls it (e.g., Millikan, 1998; Moore, 2013; Skyrms, 1996). Common ground is usually defined in terms of belief: ϕ is common ground between us iff you and I believe ϕ , each of us believes that the other believes that ϕ , each of us believes that the other believes that the other believes that ϕ , and so on (Friedell, 1969; Lewis, 1969, pp. 52–60; Stalnaker, 2002). This endless doxastic cascade has caused some scholars to fear that common ground, thus defined, must be hard if not impossible to attain, and that therefore the concept is best abandoned altogether. However, for reasons explained already by Lewis himself, such fears are unfounded. The critical point is that the doxastic cascade "is a chain of implications, not of steps in anyone's actual reasoning. Therefore there is nothing improper about its infinite length." (Lewis, 1969, p. 53; for further discussion, see Lewis, 1969, pp. 52–60; Schiffer, 1972, pp. 30–36; Vanderschraaf & Sillari, 2014.) In the following, I will argue that, without common ground, conventions just could not work, so in this point I am with Lewis.

Second, according to Lewis, conventions are solutions to coordination problems. By definition, a coordination problem is any situation that forces agents with broadly aligned interests into making interdependent decisions that may result in more than one coordination equilibrium: a combination of actions in which nobody would have been better off had any single else acted otherwise (Lewis, 1969, pp. 8–24). Lewis's list of sample coordination problems mostly consists of such mundane items as deciding where to meet, rowing a boat together, and dressing for a party. However, although it is clear that we heavily rely on language to deal with our coordination problems, it does not seem right to say that we address coordination problems whenever we use the word "bicycle," a falling intonation contour, or interrogative syntax. Therefore, as well as for reasons given by Millikan (1998), Sugden (2005), Marmor (2009), and others, I hold that the purposes of conventions are broader than Lewis makes them out to be. I do accept that conventions generally serve the purpose of action coordination, but the concept of action coordination covers a wider range of situations than what Lewis has in mind. For example, fashion in clothing is a form of action coordination which may but need not solve a coordination problem (Davis, 2003). That said, although in the following I will assume that conventions enable action coordination, my argument does not require that this be

true. All it requires is that conventions serve purposes and in particular that, occasionally, several conventions serve the same purpose; which should be uncontroversial.

Third and last, I agree with Millikan (1998), Davis (2003), Moore (2013), and others, who argue against the idea that general conformity is an essential trait of conventions. General conformity, be it *de facto*, *de jure*, or both, has often been taken to be a hallmark of conventions (e.g., Bach & Harnish, 1979; Searle, 1969), and it figures conspicuously in Lewis's analysis, too, though eventually he weakens it somewhat so as to allow for the occasional exception. However, there are plenty of cases to show that, weakened or not, general conformity is not mandatory. For example, consider some of the more common ways of saying goodbye in English: "Good bye," "Bye," "See you," "See you later / soon / tomorrow / ...," "Good night." Each of these expressions is conventionally employed to say goodbye, and although their usage conditions are not exactly the same, there is so much overlap that a speaker will rarely if ever feel obliged to pick one rather than any other.

More generally, the existence of synonyms and near-synonyms ("couch / sofa," "large / big," "sad / blue," "heart attack / myocardial infarction") belies the claim that conventions are inherently lawlike. For another example, consider the general availability of multiple media. If Donald wants to inform his mother that he passed his driving test, he can choose to tell her face to face or on the phone, he can write her a letter or email, leave a message on Twitter or Facebook, and so on. All of these are conventional means of information sharing, and in most cases no single option is preferred over all others, let alone mandatory.

Observations like these, which are easy to multiply, prove that conventions often fail to require general conformity. The essential function of conventions is to *enable* action coordination; Lewisian conventions are a special case in that they *prescribe* how a given coordination problem is to be solved. Conventions may attain lawlike status in two ways. One is that it may be in everyone's interest that a convention should be followed by all; this is the type of case that Lewis was concerned with, epitomised by the convention of driving on the right (or left, as the case may be). Another possibility, not always disjoint with the first, is that conventions are enforced, institutionally or otherwise; again, driving on the right is a case in point. I take it that the latter possibility is of secondary interest, and will confine my attention to the former.

Properly speaking, lawlike conventions are a species of enabling conventions, but in the following it will be more convenient to use the two terms contrastively. Hence, I will employ "enabling" as short for "merely enabling," thus making "lawlike" and "enabling" mutually exclusive by fiat. The distinction between lawlike and enabling conventions is an important one, because the two affect our behaviour in very different ways. Lawlike conventions as it were impose themselves on us: we stop at red and pay for groceries because we have to. By contrast, enabling conventions offer themselves to us as possible ways of regulating our social interactions, and therefore they present us with choices all the time: shall I say "hi" or "hello," "big" or "large," "sofa" or "couch"? In the following, I will argue that, whereas some types of linguistic convention are lawlike, others are merely enabling, and that the latter include sense conventions which associate lexical forms with meanings and sentence forms with speech acts. It naturally follows that, in general, sense conventions are a source of indeterminacy, which interlocutors will have to contend with.

Whether lawlike or enabling, conventions are behavioural *regularities*. Suppose we are both in London and have agreed to meet in Russell Square. This arrangement works well enough for both of us that a convention establishes itself: whenever both of us are in London we meet in Russell Square. Now at what point in our history of meetings did that convention come into existence? There is no such point: our convention emerged gradually, so it is impossible to pin down its time of birth; conventions are vague in this respect (cf. Davis, 2003; Moore, 2013). However, unless we

explicitly agree to establish Russell Square as a standard coordination device, it is clear that after our first meeting we did not have a convention going yet. One precedent does not suffice to make a convention; for that, precedents need to have occurred often enough.

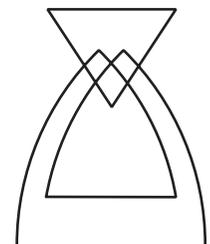
Conventions are social regularities of the past that enable us to shape our present behaviour; that is what they are for. However, although one precedent does not make a convention, one precedent may be enough to solve a problem in the present. To illustrate, let us conduct a thought experiment (cf. Schelling, 1960; pp. 67–68, Lewis, 1969, p. 36). We take two people who have never met before, introduce them to each other, and seat them in separate rooms, where they receive identical instructions:

We ask you to pick one colour from the following: green, blue, red, and yellow. The person next door is offered the same choice. If both of you pick the same colour, you earn five euros each. If you pick different colours, you get nothing.

Suppose that our subjects strike lucky: purely by chance, they both choose blue, and thus receive €5 each. Then we repeat the experiment with the same subjects. Which colour will they pick this time? It is a foregone conclusion that both will choose blue again, and this time each can be confident that the other will do so, too.

On the first trial, our subjects solved their coordination problem by sheer luck. On the second trial, they solved it on the strength of a single precedent, and in the absence of any convention, thanks to the fact that, this time, one solution was uniquely salient. Here and in the following, “salience” is a common ground concept: in order for it to enable coordination between us, a precedent must be salient for *us*, as a team, and a precedent may be salient for a team without being salient for each of its members severally. In the second trial of our thought experiment, for example, though the colour blue was salient for our two subjects as a team, one or both may have been privately attending to another colour (see Schelling, 1960 and Mehta, Starmer, & Sugden, 1994 for experimental evidence and discussion).

Single-precedent coordination is not restricted to thought experiments. From the mid-1960s onwards, Glucksberg, Krauss, and colleagues conducted a series of studies using a reference task in which participants had to identify unusual figures designed to resist straightforward description (e.g., Glucksberg, Krauss, & Weisberg, 1966; Krauss & Glucksberg, 1977). In these experiments, speaker and hearer were separated by a barrier. The speaker had before him a page showing a set of numbered figures (an example is shown on the right); the hearer had a page with the same figures, but arranged differently and without numbers. The task was for the speaker to refer to each of the figures and give its number, so that the hearer could number it accordingly. Since each figure appeared in several trials, this design made it possible to trace the emergence and development of referential expressions within a short time frame.



People turn out to be remarkably good at this task. Speakers readily produce impromptu labels like “hourglass” or “cup” to refer to the figure above, and often enough their partners immediately copy them on subsequent trials. In some cases, participants need more than one trial to agree on a label, as the following sequence of trials illustrates (Krauss & Glucksberg, 1977):

Looks like an hourglass with legs on each side ... hourglass with the legs ... hourglass-shaped thing ... hourglass ... hourglass ... hourglass.

But even so, it seldom takes more than a handful of trials for participants to converge on a label, and quite often labels are established instantaneously.

In the Glucksberg/Krauss experiments, participants were faced with coordination problems that could be solved in many different ways, as was attested by the variety of speakers' productions. However, one solution might be particularly salient because it happened to work on a previous trial, and if such a solution presented itself, it would tend to be reused. Outside the psychology lab, salient precedents are usually provided by conventions. We call a spade a "spade" because that is how spades have been called often enough, and frequency of occurrence promotes salience. But obviously, the emergence of a convention necessarily involves a run-up period, during which the salience of precedents was not yet due to their frequency.

Thus, the following picture suggests itself. We may coordinate our actions by means of precedents provided they are sufficiently salient. A linguistic device can be salient because it has occurred recently enough and with a certain minimum regularity, in which case there is a convention in place. But as we have seen, even a single occurrence may suffice to make a precedent salient. Either way, in order to understand how communicative and other social interactions of the past inform present ones, we are well advised to first investigate the kinematics of precedent-based coordination. Which is what we will do shortly; but first some preliminaries are in order.

3 | UTTERANCES

Utterances are commonly understood as actions involving the delivery of a sentence, and though it is seldom denied that productions of sub-sentential expressions are utterances, too, in general not much importance seems to be attached to this fact. However, since my objective in this paper is to analyse the conventionality of words, syntactic constructions, and so on, I need to stress that the action of uttering a complex expression never comes alone. If A tells B,

1. Theresa is fond of grass,

then in a normal run of events, A's utterance serves to inform B that Theresa is fond of grass. But part of A's utterance is an utterance of the proper name "Theresa," which helps A and B to get aligned on a goat named "Theresa," say, and likewise each of the other words is uttered singly as part of A's utterance of (1). And that is not the end of it, for in the process of uttering (1), A also utters the verb phrase "is fond of grass," the adjectival phrase "fond of grass," and the prepositional phrase "of grass."

I take it that these observations are unobjectionable, but they are more significant than is generally realised. An utterance of a sentence is an intentional action with a communicative purpose, and when the action aims to be cooperative, it will license the kind of pragmatic inferences that Grice (1975) dubbed "conversational implicatures." The exact same thing holds true of utterances of words: they are actions with communicative purposes, which are generally cooperative, and therefore give rise to Gricean inferences (though I recommend not to call such inferences "conversational implicatures," so as to forestall terminological confusion). Seen this way, communication is Gricean from the lexical level up, and the embattled boundary between pragmatics at and below the sentence level is of relatively little interest (see Geurts, 2010, pp. 182–187 for further discussion).

Bearing in mind that utterances need not be utterances of sentences also helps to see how we can approach Lewis's problem of the relation between language and languages. We can view words, quite abstractly, as pairings of forms and meanings, which in and of themselves are not

conventional. However, words may serve to individuate utterances, and since utterances are actions, and actions *can* be conventional, we can make sense of the notion that a word is conventionally used in this or that way. The same, *mutatis mutandis*, for syntactic categories and constructions, intonation contours, and so on. This idea will be developed further in the following sections.

Another aspect of utterances that bears emphasising is that, to a very large degree, they are joint projects between interlocutors.¹ It has been observed that, in spoken dialogue, speakers often take turns in mid-sentence, as in the following attested example (Clark & Wilkes-Gibbs, 1986):

2. A: That tree has, uh, uh, ...
 B: tentworms.
 A: Yeah.
 B: Yeah.

Interactions like this are common, and demonstrate the extent to which interlocutors are able to collaborate so as to achieve successful communication. In the exchange in (2), it is A who asserts that the tree under discussion has tent worms, but somewhat paradoxically, his assertion is the result of a collaborative effort; in a sense, A and B establish together what A says.

Let us develop this point a bit further. Suppose that A and B agree, wrongly, that the author of *Lolita* was Vyacheslav Molotov. Now A says:

3. The author of *Lolita* was born in 1890.

Since it is common ground between A and B that Molotov was the author of *Lolita*, A's utterance means for them that Molotov was born in 1890 (which happens to be true). As far as A and B are concerned, "the author of *Lolita*" successfully refers to Molotov, and A's utterance of (3) successfully informs B that Molotov was born in 1890. That A and B are wrong about the facts is neither here nor there.

Now A finds that *Lolita* wasn't written by Molotov after all, but sadly he still gets things wrong: for reasons we need not go into here, he has come to believe that Sergei Prokofiev (b. 1891) is the author, and he corrects his earlier statement accordingly:

4. The author of *Lolita* was born in 1891.

While communication was successful in the first case, this time something has gone seriously awry: A intends to refer to Prokofiev and B takes him to be referring to Molotov. To make matters worse, A *knows* that B will take him to be referring to Molotov. A and B fail to agree on the reference of "the author of *Lolita*," and therefore communication has failed.

These two cases illustrate that reference is not determined by the speaker alone but by what is common ground between interlocutors. In the first case, it was common ground between A and B that Molotov wrote *Lolita*, and reference was successful. In the second case, reference failed because they disagreed about the authorship and only A was aware that they did.² The same holds, *mutatis mutandis*, for meaning. If it is common ground between A and B that "molybdenum" and "wolfram"

¹ In this and the following paragraphs, I will be merely scratching the surface of a theme discussed at much greater length by Clark (1992, 1996).

² Some philosophers would claim that, as a matter of fact, the referent of "the author of *Lolita*" is the same in both cases, viz. Vladimir Nabokov. This view strikes me as dubious, at the best of times (cf. Geurts, 1997), but that is as it may be, since I am concerned with communication rather than the metaphysics of reference.

are synonymous, then the fact that their usage disagrees with the experts' will not prevent them from communicating successfully about what they take to be a single metal.

The upshot of these observations is that, at the end of the day, reference and meaning are determined neither by the speaker's intentions nor by the conventions of the interlocutors' national language, but by their common beliefs about the worldly and linguistic facts. It seems likely that interlocutors' common beliefs about language and the world will predominantly agree with their private beliefs and the beliefs of the larger communities they belong to. But wherever they disagree, it is the common ground that prevails. Precedents and conventions outside the common ground are communicatively inert.

4 | PRECEDENTS

With these preliminaries out of the way, let us turn to precedents. The basic idea is simple enough: if it is common ground between interlocutors that a given expression or construction α has previously been used (by themselves or others) for a particular purpose p , then in principle α may be reused for the same purpose. Precedents are past actions that will serve as templates for present ones. In principle, it does not matter how often α has been used to achieve p , though popular precedents are more likely to be reused. Even so, as we have seen, one precedent may be enough, provided it is sufficiently salient.

I distinguish between two types of utterances: speech acts and alignments. I will not try to define what speech acts are, since for my current purposes it is not necessary to take a stance on that issue (but see Geurts, 2017, *In press*). All I need to assume here is that utterances may serve to inform, request, ask, and so on, that it is an essential part of linguistic communication to establish what speech acts are being performed, and that a single utterance may serve to perform several speech acts at once; an utterance of "It's cold outside," for example, may count as a warning and a reminder at the same time. None of these assumptions should be particularly controversial.

If speech acts are the princesses of linguistic interaction, alignments are their handmaidens. Their purpose is joint attention. If I utter the name "John," for example, then I am trying to get us aligned on a particular individual I have in mind. More generally, an expression α serves to align speaker and hearer on an entity of α 's semantic type. Thus, if our semantic theory of choice stipulates that declarative sentences express propositions, then an utterance of "It's cold outside" will usually serve to align the interlocutors on the proposition that it is cold outside. ("Usually," because an utterance may fail to get the interlocutors aligned on anything.) I assume that, with respect to alignments, monosemy is the norm; that is to say, no utterance may serve to align interlocutors on more than one semantic object at a time. Therefore, an utterance of "It's cold outside" usually serves to align the interlocutors on a single proposition. However, note that this does not rule out the possibility that the same utterance serves to perform a speech act, or several. For example, "It's cold outside" typically expresses the proposition that it is cold outside and may at the same time be a statement, a reminder, and a warning.

To make these ideas more concrete, let us introduce the following toy syntax. "John snores" is parsed as $S:[N:John, V:snores]$, where S is the category of declarative sentences, and N and V subsume proper names and intransitive verbs, respectively. Small capitals are used to represent semantic entities: $JOHN_{43}$ and $JOHN_{76}$ are distinct individuals; $SNORE$ is a concept, which combines with an individual to yield a proposition (e.g., $SNORE(JOHN_{43})$); and $REMIN$ and $WARN$ are speech acts, which we model as relations between speakers, hearers, and propositions. Last, we introduce the following shorthand notation:

$P_t(\alpha, x) : =$ at t an utterance of α enabled interlocutors to align on x , where t is a time in the past.

Hence, $\exists t P_t(\text{N:John}, \text{JOHN}_{43})$ is short for: “Some time ago, an utterance of the name “John” enabled interlocutors to align on JOHN_{43} .”

Now we can give a very simple model to illustrate how precedents interact so as to enable communication between two individuals A and B:

5. It is common ground between A and B that:
 - a. $\exists t P_t(\text{N:John}, \text{JOHN}_{43})$.
 - b. $\exists t P_t(\text{N:John}, \text{JOHN}_{76})$.
 - c. $\exists t P_t(\text{V:snores}, \text{SNORE})$.
 - d. $\exists t P_t(\text{S}:[\text{N}:\alpha, \text{V}:\beta], c(x))$, for some individual x and concept c such that $P_t(\text{N}:\alpha, x)$ and $P_{t'}(\text{V}:\beta, c)$, for some t', t'' prior to t .
 - e. $\exists t$: S:[...] served to REMIND at t that p , where $P_t(\text{S}:[...], p)$.
 - f. $\exists t$: S:[...] served to WARN at t that p , where $P_t(\text{S}:[...], p)$.

Each of the clauses in (5) states that there is at least one precedent for using a given word or construction in a certain way. All this is common ground between A and B, and to keep things as simple as possible we assume that there are no further precedents in the common ground; that is to say, it is common belief between A and B that there are only two types of precedent for utterances of “John,” only one type for utterances of “snores,” and so on.

In contrast to the lexical precedents in (5a–c), the precedent in (5d) involves a grammatical construction. (5d) says that, at some t in the past, a sentence S occurred consisting of a name α and a verb β ; that α and β had served, before t , to align interlocutors on an individual x and a concept c , respectively; and that S served at t to align interlocutors on the proposition $c(x)$. Note that this precedent overlaps with the ones in (5e) and (5f), which specify that declarative sentences have been used to issue a reminder and a warning, respectively.

Now suppose A utters “John snores,” and suppose furthermore that the interpretation of A’s utterance is constrained solely by the precedents in (5), and that B is entitled to assume that A is performing a speech act. Then B is presented with two choices: she must decide (i) whether the relevant precedent for “John” is (5a) or (5b), and (ii) which of (5e) and (5f) are the precedents for A’s uttering “John snores” (recall that an utterance may count as several speech acts at once). Hence, there are six possible readings for A’s utterance of “John snores”; that is to say, six ways of linking that utterance and its constituent utterances to the precedents in (5):

- REMIND(A, B, SNORE(JOHN₄₃))
- REMIND(A, B, SNORE (JOHN₇₆))
- WARN(A, B, SNORE(JOHN₄₃))
- WARN(A, B, SNORE (JOHN₇₆))
- REMIND(A, B, SNORE(JOHN₄₃)) and WARN(A, B, SNORE(JOHN₄₃))
- REMIND(A, B, SNORE(JOHN₇₆)) and WARN(A, B, SNORE(JOHN₇₆))

Note that the precedents in (5c) and (5d) are implicated in each of these readings.

The model in (5) was designed to capture, on a small scale, some of the ambiguities that are endemic to natural language: as a rule, a name has many bearers and any grammatical mood can be

used to perform many different speech acts. These and other types of ambiguity combine to generate a multitude of candidate readings for even the simplest utterance. To a large degree, this explosion is contained by the fact that precedents outside the common ground can be discarded. For example, although it is a fact, and A and B may be aware, that there are millions of Johns in this world, only two of them are in their common ground, which radically reduces the indeterminacy of A's uttering "John."

Nonetheless, ambiguities remain, and even in the simplest cases they add up, as our example shows. Fortunately, we manage to resolve these ambiguities, more often than not, homing in on that reading which stands out as the most likely one, given the discourse purposes and other contextual factors. Here, too, common ground plays a key role, for in general, there will be mutual agreement on what is the best way of resolving a given ambiguity. For example, if JOHN₄₃ is the current topic of conversation, and A utters the name "John," then it will be common ground between A and B that JOHN₄₃ is the most likely referent. Therefore, B is entitled to assume that A's utterance targeted JOHN₄₃, it is common ground that she is so entitled, and by the same token, A is committed to having targeted JOHN₄₃, regardless what his referential intentions were. For another example, suppose it is common ground that B did not know that John snores. Then within the scenario of (5), it is common ground that A's utterance of "John snores" cannot be a reminder. Therefore, it is common ground that B is entitled to suppose that A's utterance serves to warn, not to remind, and by the same token, A is committed to having issued a warning, not a reminder.

Thus, the common ground enables the use of precedents in at least two ways. First, it restricts the range of precedents eligible for speakers to emulate, without which the number of candidate precedents would generally be too large to be manageable. Second, the common ground allows interlocutors to negotiate their way out of any remaining ambiguities, by supporting abductive inferences about the preferred ways of resolving indeterminacies. Without common ground, precedent-based interaction would be unthinkable.

5 | FROM PRECEDENTS TO CONVENTIONS

The model in (5) illustrates how words and constructions help to individuate and guide the combinatorics of the precedents underpinning linguistic communication. Conventions are not part of the picture yet, but it will not be hard to see where they come in. As things currently stand, neither the recency of a precedent nor the frequency with which it has occurred can affect the likelihood that it will be copied. Therefore, a single deviant use of a word in the distant past will continue to compete with what has long become its dominant use. Clearly, we must keep track of when and how often a given *type* of precedent has occurred, and I take it that this is what speakers do. So we need precedent types, which we may think of as sets of precedents. (This is admittedly crude, but it will do.) Conventions can then be defined as a special kind of precedent type:

6. If R is a precedent type all of whose instances served the same purpose *p*, then R is a convention in a population P iff R has been instantiated often enough and recently enough that it is a salient way of achieving *p* for the members of P.

Obviously, this is a sketchy definition at best, if only because it contains several vague terms; but that is the nature of the beast. Precedence is the key notion; conventions merely occupy a vague region in precedent space. The main idea underlying (6) is simply that the salience of a precedent

type varies with the recency and frequency of its instances, and the more salient a precedent type is, the more strongly it suggests itself as a coordination device.

To illustrate this idea, let P_J be a precedent type corresponding to (5a), that is, a set of past utterances of the proper name “John” that enabled interlocutors to align on $JOHN_{43}$. Suppose that P_J contains sufficiently many utterances that are recent enough to render it salient to A and B that uttering “John” is a way of aligning on $JOHN_{43}$. Then according to (6), uttering “John” is a conventional method, between A and B, for aligning on $JOHN_{43}$. In the same way, the remaining clauses in (5) map to candidate conventions between A and B. I call such conventions “sense conventions,” and divide them, along the same line as precedents, into alignment conventions and speech act conventions.

As defined in (6), it is sufficient for conventionhood that a precedent type enables a course of action. Lawlike conventions of the kind analysed by Lewis are special cases, which emerge when it is in the interest of all to adopt the same pattern of behaviour. Sense conventions are typically (merely) enabling. By contrast, many language conventions are Lewisian. In English, articles must sit at the left periphery of the noun phrase, subjects and finite verbs must agree in person and number, a syllable onset must have at most three consonants, and so on. Such form conventions are quite different from sense conventions, and it is not entirely obvious what they are for. I conjecture that their main purpose is merely to curtail the vast number of forms speakers and hearers would otherwise have to cope with. For example, if any old combination of English phonemes counted as a syllable of English, the number of potential syllables would be prohibitively large. Thankfully, phonotactic conventions reduce that number dramatically, thus making it humanly possible to produce and parse syllables.³

If this conjecture is on the right track, form conventions may be seen as coordination devices, too, but whereas the purpose of sense conventions is to enable interlocutors to coordinate on content, the purpose of form conventions is to coordinate on a shared pool of eligible expressions. This is why form conventions tend to require general conformity. Form conventions are metaconventions in the sense that, by curtailing admissible forms, they constrain the range of possible sense conventions.

6 | SENSE CONVENTIONS

To return to sense conventions, let us compare and contrast the two main varieties distinguished in the foregoing. Traditionally, the study of (what I call) alignment conventions has been segregated from the study of speech act conventions, which may have fostered the impression that they have little if anything in common. In this section, I argue, however, that the two types of conventions are similar, and that even if speech acts and alignments are very different types of action, the ways in which they are negotiated are much alike. In order to bring home this point I will focus on the parallels between lexical alignments and speech acts, following Asher and Lascarides (2001).

To begin with, in lexical alignments as well as in speech acts, content is underdetermined by form. Lexical indeterminacies are positively virulent. Practically, all words in everyday use have more than one sense that is common ground throughout large parts of the language community. The same holds for grammatical moods, prosody, and other “force-indicating devices,” as Searle (1969) calls them: on any given occasion, the linguistic form of an utterance constrains but does not determine its speech act(s), just as the form of a word constrains but does not determine its contextual meaning. The mere fact that a handful of grammatical moods suffice to perform an indefinite variety of speech acts betrays the extent of the indeterminacy.

³ To make this more concrete: if the 40-odd phonemes of English were allowed to combine freely, we would get 2^{40} biphonemic combinations, 3^{40} triphonemic combinations, and so on. These are frighteningly large numbers. Phonotactic constraints leave something on the order of 15,000 syllables in actual use.

Second, the various senses of a word are typically related. It is no accident that the word “table,” for example, may refer to a piece of furniture with a flat top, a group sharing a table, various kinds of flat surfaces, and so on. Similarly, the word “school” has a variety of conventional senses interconnected by world knowledge:

Our school is ...
 ... on holiday. (staff/students)
 ... on the list of institutions currently being reviewed by the Department of Education.
 (institution)
 ... on fire. (building)

In like manner, a declarative sentence may be used to issue a variety of interrelated speech acts: statements, promises, threats, warnings, complaints, and so on. Hence, the declarative mood is polysemous in much the same way as the bulk of the lexicon is.

It is important to note, *pace* Pustejovsky (1995) and others, that even if polysemy is in many ways a regular phenomenon, the semantic lexicon is not a purely generative system; for derived senses may have to be sanctioned by convention (Briscoe & Copestake, 1999; Copestake, 2012; Nunberg, 2004). Nunberg lists several cases that illustrate this point: in French, but not in English, names of fruits may be used to refer to the brandies made from them (“une poire,” “une prune”); in Russian, but not in English, the name of an organ may denote a disease of that organ; and while in English, the same noun may generally be used for an animal and its meat, in Greenlandic Eskimo this is not permitted, rendering the equivalent of “I don’t eat walrus” infelicitous. Lexical meaning shifts are enabled by world knowledge: from an established word sense, a related sense that is sufficiently salient may be derived. But still, the derivation may have to be licensed by convention.

The same is true of speech acts. Like lexical meanings, speech acts perforce require linguistic vehicles to be conveyed: grammatical constructions, particles, prosodic patterns, and so on; and there are conventional restrictions on what form is suitable for what purpose. A request like “Would you be so kind to peel the potatoes?” is generally an acceptable way of inviting the addressee to peel the potatoes, but it would be out of place in a recipe, even if it might be perfectly intelligible. When it comes to giving instructions, bare imperatives are the norm. Similarly, while the Dutch transliteration of “You are Agnes or not?” is a licit tag question, in English it is not, though the speaker’s meaning might be clear enough. As in the case of lexical polysemy, even if the intended meaning of a linguistic vehicle is fully transparent, it may have to be sanctioned by convention.

Once a derived lexical sense has conventionalised, it often remains transparent that it is a derived sense. For example, though it is fully conventional to measure lengths in “feet,” the mensural sense of the word “foot” is clearly derived from its anatomical sense. The relation between an indirect speech act and its source is similar. As has often been observed, the requestive use of “Could you pass me the salt?” obviously derives from its interrogative use, and though both uses are conventional, the derivational link remains transparent. In fact, the same form can function as a question and a request at the same time, as is shown by the fact that the addressee may respond by saying, “Yes, of course,” while giving the salt, thus answering the question and heeding the request at the same time. Given our assumption that one utterance may count as several speech acts at once, this is to be expected. By contrast, having assumed that, for alignments, monosemy is the norm, we predict that it is unusual for a word to have two senses at the same time, which, barring puns, seems to be correct.

As described by Millikan (1998, p. 176), a language is “a tangled jungle of overlapping, criss-crossing traditional patterns,” with ambiguities springing up where patterns overlap. The picture is

apt, and it highlights the importance of common ground, for it is hard to see how we could find our way around Millikan's jungle of multifarious and interrelated meanings without a shared basis enabling us to align on the same meaning most of the time. However, being at pains to play up the continuity between us and the rest of the animal kingdom, Millikan herself chooses to eradicate all "intellectualist" notions from her account of human communication, and that includes common ground (cf. also Skyrms, 1996). This leaves her theory with serious explanatory gaps. Most strikingly (and ironically), it fails to explain how hearers manage to cope with the rampant ambiguities caused by the very "crisscrossing conventions" that Millikan herself lays so much emphasis on. All she has to say on this matter is that hearers sort out ambiguities "by one means or another" (1998, p. 176), and without any notion of common ground, there is little more for her to say.

On my account, conventions and common ground intersect in various ways, one of which is especially important in the present context: the common ground constrains the range of conventions that might be relevant for the ongoing interaction. Onomastic conventions provide a particularly vivid case in point. As noted before, there are millions of individuals called "John"; that is several million of sense conventions which associate that name with an specific boy, man, baboon, or what have you. The remarkable fact that my uttering "John" enables us to align on the same individual can only be explained by supposing that we winnow down that vast set to those individuals in our common ground who go by the name of "John," and then pick out whomever we agree to be the best candidate, that is, the best candidate according to our common ground. And so it goes with sense conventions in general. To reject the notion of common ground is to give up on explaining even the most basic forms of communication.

Following Lewis, I have taken the view that conventions are, first and foremost, regularities in social interaction. This leaves out of account the fact that conventions seem to be normative in some sense (not to be confused with the notion that *speech acts* may be normative in some sense; see Kissine, 2013, pp. 158–161, Ball, 2014). But how could a mere regularity be normative? The discussion of the last paragraph suggests an answer. When I utter the word "table," then in many cases you thereby become entitled to suppose that I am trying to get us aligned on the TABLE concept, that is to say, the concept of a piece of furniture with a flat surface for eating, writing, and so on; and by the same token, I become committed to trying to get us thus aligned. However, as I have been stressing, the same word is involved in other sense conventions, too. The convention of using "table" to mean TABLE *becomes* normative in some contexts only. These are contexts in which it is common ground between us that (i) this convention is available, and (ii) this use allows us to make better sense of my utterance than any other for which the common ground offers precedents (conventional or not). If, and only if, these conditions are met, it is common ground between us that you are entitled to interpret my utterance of "table" as a bid to align on the TABLE concept, and my commitment is simply the obverse of your entitlement. In the same way, my uttering, "Does your dog bite?" will typically (though not invariably) commit me to having a wish for information, and entitle you to assume that I have that wish.

On this view, sense conventions are not intrinsically normative. What normativity they seem to have depends entirely on the context, consisting as it does of the commitments and entitlements flowing from the common ground via abduction.

7 | CONCLUSION

People are social animals in constant need of having to coordinate their actions. To a very large degree, they manage to do so by reusing solutions that worked before, for themselves or others. Like

all animals, we are creatures of habit, but the sheer variety of our social interactions force us to be more flexible than most. Each of us enters the social arena with a stockpile of memories of past interactions, clustered into criss-crossing types and subtypes; these are the precedents we draw upon for coordinating our actions. In essence, this is Lewis's picture, which I have developed in two main ways: by dropping his stipulation that conventions require general conformity, and by drawing on the notion of common ground even more heavily than Lewis did. Thus, it became possible to account for a seemingly simple fact that eluded Lewis's analysis, namely, that words, constructions, intonation contours, and other linguistic devices are conventional means for coordinating our communicative interactions.

ACKNOWLEDGEMENTS

I would like to thank Teun Dupont, Mikhail Kissine, Richard Moore, and Marat Shardingaliev for their helpful comments on previous versions of this paper.

REFERENCES

- Asher, N. & Lascarides, A. (2001). Indirect speech acts. *Synthese*, 128, 183–228.
- Bach, K. & Hamish, R. M. (1979). *Linguistic communication and speech acts*. Cambridge, MA: MIT Press.
- Ball, B. (2014). Speech acts: Natural or normative kinds? The case of assertion. *Mind & Language*, 29, 336–350.
- Briscoe, E. J. & Copestake, A. (1999). Lexical rules in constraint-based grammars. *Computational Linguistics*, 25, 487–526.
- Clark, H. H. (1992). *Arenas of language use*. Chicago: University of Chicago Press.
- Clark, H. H. (1996). *Using language*. Cambridge: Cambridge University Press.
- Clark, H. H. & Wilkes-Gibbs, D. (1986). Referring as a collaborative process. *Cognition*, 22, 1–39.
- Copestake, A. (2012). The semi-generative lexicon: Limits on lexical productivity. In J. Pustejovsky, P. Bouillon, H. Isahara, K. Kanzaki & C. Lee (Eds.), *Advances in generative lexicon theory* (pp. 455–474). Berlin: Springer.
- Davis, W. A. (2003). *Meaning, expression, and thought*. Cambridge: Cambridge University Press.
- Friedell, M. F. (1969). On the structure of shared awareness. *Behavioral Science*, 14, 28–39.
- Geurts, B. (1997). Good news about the description theory of names. *Journal of Semantics*, 14, 319–348.
- Geurts, B. (2010). *Quantity implicatures*. Cambridge: Cambridge University Press.
- Geurts, B. (2017). Fictional commitments. *Theoretical Linguistics*, 43, 53–60.
- Geurts, B. (In press). Making sense of self talk. *Review of Philosophy and Psychology*.
- Glucksberg, S., Krauss, R. M. & Weisberg, R. (1966). Referential communication in nursery school children: Method and some preliminary findings. *Journal of Experimental Child Psychology*, 3, 333–342.
- Grice, P. (1975). Logic and conversation. In P. Cole & J. Morgan (Eds.), *Syntax and semantics 3: Speech acts* (pp. 41–58). New York, NY: Academic Press.
- Hawthorne, J. (1990). A note on 'languages and language'. *Australasian Journal of Philosophy*, 68, 116–118.
- Kissine, M. (2013). *From utterances to speech acts*. Cambridge: Cambridge University Press.
- Krauss, R. M. & Glucksberg, S. (1977). Social and nonsocial speech. *Scientific American*, 236(2), 100–105.
- Lewis, D. K. (1969). *Convention*. Cambridge, MA: Harvard University Press.
- Lewis, D. K. (1975). Languages and language. In *Minnesota studies in the philosophy of science* (pp. 3–35). Minneapolis: Minnesota University Press.
- Marmor, A. (2009). *Social conventions*. Princeton: Princeton University Press.
- Mehta, J., Starmer, C. & Sugden, R. (1994). The nature of salience: An experimental investigation of pure coordination games. *The American Economic Review*, 84, 658–673.
- Millikan, R. G. (1998). Language conventions made simple. *The Journal of Philosophy*, 95, 161–180.
- Moore, R. (2013). Imitation and conventional communication. *Biology and Philosophy*, 28, 481–500.
- Nunberg, G. D. (2004). The pragmatics of deferred interpretation. In L. R. Horn & G. Ward (Eds.), *The handbook of pragmatics* (pp. 344–364). Oxford: Blackwell.
- Pustejovsky, J. (1995). *The generative lexicon*. Cambridge, MA: MIT Press.
- Rescorla, M. (2014). Convention. In *Stanford Encyclopedia of Philosophy*. Retrieved from <https://plato.stanford.edu/entries/convention/>
- Schelling, T. C. (1960). *The strategy of conflict*. Cambridge, MA: Harvard University Press.
- Schiffer, S. (1972). *Meaning*. Oxford: Oxford University Press.
- Searle, J. R. (1969). *Speech acts: An essay in the philosophy of language*. Cambridge: Cambridge University Press.
- Skyrms, B. (1996). *Evolution of the social contract*. Cambridge: Cambridge University Press.

- Stalnaker, R. C. (2002). Common ground. *Linguistics and Philosophy*, 25, 701–721.
- Sugden, R. (2005). *The economics of rights, co-operation, and welfare* (2nd ed.). Basingstoke: Palgrave Macmillan.
- Vanderschraaf, P. & Sillari, G. (2014). Common knowledge. In *Stanford Encyclopedia of Philosophy*. Retrieved from <https://plato.stanford.edu/entries/common-knowledge/>

How to cite this article: Geurts B. Convention and common ground. *Mind Lang*. 2018;33: 115–129. <https://doi.org/10.1111/mila.12171>