University of Reading

Must we measure what we mean?

Article

Accepted Version

Hansen, N. ORCID: https://orcid.org/0000-0001-5074-1075
(2017) Must we measure what we mean? Inquiry, 60 (8). pp. 785-815. ISSN 1502-3923 doi: 10.1080/0020174X.2017.1310669 Available at https://centaur.reading.ac.uk/70011/

It is advisable to refer to the publisher's version if you intend to cite from the work. See Guidance on citing.

To link to this article DOI: http://dx.doi.org/10.1080/0020174X.2017.1310669

Publisher: Taylor & Francis

All outputs in CentAUR are protected by Intellectual Property Rights law, including copyright law. Copyright and IPR is retained by the creators or other copyright holders. Terms and conditions for use of this material are defined in the End User Agreement.

www.reading.ac.uk/centaur

CentAUR

Central Archive at the University of Reading

Reading's research outputs online

Must We Measure What We Mean?*

Nat Hansen

Forthcoming in *Inquiry*

Abstract

This paper excavates a debate concerning the claims of ordinary language philosophers that took place during the middle of the last century. The debate centers on the status of statements about "what we say". On one side of the debate, critics of ordinary language philosophy argued that statements about "what we say" should be evaluated as empirical observations about how people do in fact speak, on a par with claims made in the language sciences. By that standard, ordinary language philosophers were not entitled to the claims that they made about what we would say about various topics. On the other side of the debate, defenders of the methods of ordinary language philosophy sought to explain how philosophers can be entitled to statements about what we would say without engaging in extensive observations of how people do in fact use language. In this paper I defend the idea that entitlement to claims about what we say can be had in a way that doesn't require empirical observation, and I argue that ordinary language philosophers are (at least sometimes) engaged in a different project than linguists or empirically minded philosophers of language, which is subject to different conditions of success.

1 The historical scene of the debate

In 1957, there was an official confrontation between philosophers with radically opposed estimations of the significance of ordinary language philosophy. On one side was Stanley Cavell, then still working on his Harvard dissertation but already teaching at UC Berkeley. Cavell had been deeply impressed by Austin's practice of ordinary language philosophy while attending his 1955 classes at Harvard (Cavell 2010, pp. 322–326), and he lectured on Austin in the spring term of 1957 at Berkeley (pp. 334–335).[1] On the other side of the

*This project is supported by an External Faculty Fellowship at Stanford's Humanities Center and a Leverhulme Research Project Grant, RPG-2016-193. Thanks to Eliot Michaelson, Chauncey Maher, Dan Lassiter, Jay Elliott, Bruno Ambroise, Zed Adams, participants in the Cavell seminar at the New School for Social Research and the fellows' workshop at Stanford University's Humanities Center for very helpful comments.

[1] Austin himself was at Berkeley in 1958–1959, giving the series of lectures that would become *Sense and Sensibilia*—see Cavell (1976b, p. 97 n. 3).

confrontation was his colleague Benson Mates, a logician and historian of ancient philosophy. Cavell says, "My insistence on the treasures I was finding Austin to have brought to philosophy was getting on the nerves of some accomplished teachers in and around my senior colleagues in the Berkeley department", and he was invited (or "ordered" (p. 360)) to defend the methods of ordinary language philosophy at the 1957 Christmas Pacific APA at Stanford. His opponent was Mates, who himself taught a seminar on Austin's "A Plea for Excuses" in 1957. Mates's seminar was attended by several visiting advocates of "empirical semantics", a midcentury version of experimental philosophy that originated in Norway. The empirical semanticists rejected "anti-empirically oriented armchair [philosophy]" (Naess 1938, p. viii; quoted in Murphy 2015, p. 325) and conducted surveys of how philosophically significant expressions were used by non-philosophers. Mates and the empirical semanticists in his seminar were skeptical about Austin's informal methods of investigating ordinary language, and Mates's contribution to the APA symposium grew out of notes from his seminar (Murphy, 2015, p. 350–352).

Cavell reports that his response to Mates at the APA was "well enough received", meaning "that those who were pulling for me thought I had answered Mates's objections, and those who were not pulling for me thought I had lived to fight another day" (Cavell, 2010, p. 361). Mates and Cavell published their contributions to the symposium in the first volume of the new Norwegian journal *Inquiry*, founded by Arne Naess, which also went on to publish experimental work by the empirical semanticists (see Chapman 2011).

Cavell's defense of the methods of ordinary language philosophy against Mates's criticisms, "Must We Mean What We Say?", was, Cavell says, "the first text I published that I still care about" (Cavell 2010, p. 361), and Mulhall (1999, p. 1) says that "Cavell's response to Mates underpins everything else [Cavell] has written". Cavell also reports that, according to Bernard Williams, Austin had "pushed to have graduate students and younger dons at Oxford" read it (Cavell 2010, p. 149). But it also attracted intense criticism from advocates of the new, scientifically-minded philosophy of language being developed at MIT (Fodor and Katz, 1963). Cavell, recalling the message of Fodor and Katz (1963), says they asserted

> ...(I believe I can recall the exact words) that the articles ["Must We Mean What We Say" and "The Availability of Wittgenstein's Later Philosophy"] were 'deleterious to the future of philosophy' (p. 442).

Cavell is slightly misremembering here. What Fodor and Katz actually say is:

> The position Cavell advocates in [the two articles] seems to us, however, to be mistaken in every significant respect and to be pernicious both for an adequate understanding of ordinary language philosophy and for an adequate understanding of ordinary language (Fodor and Katz, 1963, p. 57).

There were later defenses of Cavell (Henson, 1965; Bates and Cohen, 1972), and attempts to defend the methods of ordinary language philosophy against Mates and Fodor and Katz, whether or not such defenses would be endorsed by Cavell (Searle, 1969; Vendler, 1967). And the Mates-Cavell debate was reprinted in two volumes focusing on the methods

of ordinary language philosophy (Chappell, 1964; Lyas, 1971). But the debate eventually fizzled out for whatever reasons philosophical debates end even when the philosophical problems motivating the debate haven't been resolved (lack of interest, change of fashion). Recently, however, questions about the proper methodology of philosophy of language and the relevance of formal experiments to the investigation of meaning have been posed again by 21st century advocates of experimental philosophy and their critics, and by new advocates of ordinary language philosophy.[2] Reexamining the 1957 debate permits an assessment of the success or failure of the original arguments for and against ordinary language philosophy from a sufficiently long historical distance to provide a new angle on the old debate.

2 Mates's challenge

Mates wonders "how one would go about verifying" the statements made about language by "the so-called ordinary language philosophers" (Mates, 1958, p. 161). He takes as his target Ryle's remarks on "the ordinary sense of 'voluntary', [and] 'involuntary'" (Ryle, 1949, p. 59).

Ryle's discussion takes place in the context of his attack on "The Will" and the idea of "volitions" understood as "special acts, or operations 'in the mind'... which somehow puts my muscles into action", which are part of the "myth of the ghost in the machine" (p. 50). Ryle objects to the idea that intentional action involves a mental "volition" partly on the grounds that "no one, save to endorse the theory, ever describes his own conduct, or that of his acquaintances, in the recommended idioms" (p. 51). He goes on to describe the way that "ordinary folk" apply the words "voluntary" and "involuntary", in contrast to the "stretched" sense that philosophers who propose a theory of volitions do, in which an action is voluntary just in case it is caused by a volition, and involuntary just in case it is not. Here is how Ryle characterizes the ordinary use of "voluntary" and "involuntary":

> In their most ordinary employment 'voluntary' and 'involuntary' are used, with a few minor elasticities, as adjectives applying to actions which ought not to be done. We discuss whether someone's action was voluntary or not only when the action seems to have been his fault...
>
> In this ordinary use, then, it is absurd to discuss whether satisfactory, correct or admirable performances are voluntary or involuntary (p. 56).

Philosophers "stretch" the ordinary sense of "voluntary" and "involuntary" when they maintain that those expressions are jointly exhaustive descriptions of human action, which leads to a "tangle of largely spurious problems, known as the problem of the Freedom of the Will" (p. 57). In order to avoid that problem, according to Ryle, we need to remind

[2]For examples of neo-ordinary language philosophy see Baz (2012), DeRose (2005) Fischer (2014), and Laugier (2013). Hansen (2014) surveys varieties of contemporary ordinary language philosophy. For comparisons of the methods of experimental philosophy and ordinary language philosophy, see Sandis (2010) and Hansen and Chemla (2015). For historical context with a focus on "empirical semantics", see Murphy (2015).

ourselves of the ordinary uses of "voluntary", "involuntary" and related expressions, and see how they do not require positing any unverifiable inner acts of will.[3]

How, Mates wonders, would we go about verifying whether Ryle is right about the way "voluntary" and "involuntary" are ordinarily used? The most obvious approach would be to look at how those expressions are in fact used. Mates is doubtful that the "ordinary man" only applies "voluntary" and "involuntary" to actions that ought not to be done, but he doesn't undertake such a study (Mates, 1958, p. 163). He points out that even if it were the case that the expressions are only applied to actions that ought not to be done, that wouldn't establish that doing so was an essential part of the meaning of those expressions—it might be the case, for example, that "some other factor, such as perhaps a disposition on the part of the ordinary man to talk more about things of which he disapproves than about things of which he approves", could account for the fact (if it were a fact) that "voluntary" and "involuntary" are ordinarily used as Ryle describes.[4]

But Ryle thinks that the obvious approach to verifying statements about ordinary use is "philosophically pointless". He distinguishes what he calls the "use" of an expression and the "usage" of that expression:

> Usage... can be local or widespread, obsolete or current, rural or urban, vulgar or academic. There cannot be a misusage any more than there can be a miscustom or a misvogue. The methods of discovering usages are the methods of philologists (Ryle, 1953, p. 174–175).

(... philologists, that is, and *not philosophers*.)

"Use", in contrast, is a "way of operating with... a word", a "technique, knack, or method". It is possible to misuse an expression, indeed, it might be the case that popular usage of an expression substantially diverges from the use of that expression. The "appeal to prevalence [of usage] is philosophically pointless, besides being philosophically risky" (Ryle 1953, p. 177). Ryle therefore says he is interested in use, not usage.[5] Mates glosses Ryle's methodology of "use" over "usage" as indicating that there is something "normative" about ordinary use, where he understands "normative" to mean that a recommendation is being made, or that a "use is sanctioned by authority" (p. 165):

> [Ryle's discussion] appears to indicate that for him there is some sort of normative element in assertions about ordinary use. If the opposite of use is misuse, then use must be somehow right, proper, or correct (Mates, 1958, p. 164).

[3] Austin (1962) argues in a similar way that philosophers have "stretched" the meaning of "directly" in the phrase "directly perceive" to the point of making the expression meaningless.

[4] This worry is an early example of what has become a standard response to the arguments of the ordinary language philosophers—offering a pragmatic explanation of a purportedly semantic phenomenon. See Grice (1961) for another early example of this type of response.

[5] Contrast Ryle's view with Horace's in *Ars Poetica*: "Many words will be born again that have now sunk into oblivion, and many will die that are now held in respect, if that's what usage chooses—usage, which has the power over the judgment, the law, and the rule of speech" (translation from Lynch 2009, p. 34).

But Mates dismisses the idea that there is a normative element in assertions about ordinary use for ordinary language philosophers, even though Ryle says that there is: "Despite Ryle's own explanations, I am reluctant to believe that the expression 'ordinary use' is really a normative term for the ordinary-language philosophers". Instead, Mates proposes that "ordinary use" is a "rough descriptive term" (Mates, 1958, p. 164) where the standards for evaluating statements about ordinary use are those appropriate for evaluating any other descriptive statement, including statements about, e.g., whether some particular usage (in Ryle's terms) of an expression occurs more frequently than some other usage. And Mates argues that the "armchair" methods of the ordinary language philosophers do not entitle them to claim that language is ordinarily used in the way they claim it is ordinarily used, on the grounds that "even relatively careful authors are often not reliable reporters of their *own* linguistic behavior, let alone that of others" and "the intuitive findings of different people, even of different experts, are often inconsistent" (Mates, 1958, p. 165).

As an example of expert disagreement about ordinary language, Mates cites Austin discussing the expressions "involuntary/involuntarily" in a way that seems to contradict Ryle's observations about the ordinary use of "voluntary" and "involuntary", as applying only to actions which ought not to be done:

> . . . for example, take 'voluntarily' and 'involuntarily': we may join the army or make a gift voluntarily, we may hiccough or make a small gesture involuntarily. . . (Austin, 1957, p. 17).

So there is an apparent disagreement between Ryle and Austin, and, as Mates puts it, "If agreement about usage cannot be reached within so restricted a sample as the class of Oxford Professors of Philosophy, what are the prospects when the sample is enlarged?" (Mates, 1958, p. 165).[6]

Mates also makes a positive proposal for how one should go about verifying statements about the ordinary use of expressions. He discusses two complementary approaches to verifying descriptive statements about how words are used: the "extensional" approach, and the "intensional" approach.

> In the extensional approach one observes a reasonably large class of cases in which the subject applies the word, and then one "sees" or "elicits" the meaning by finding what is common to these cases . . . In the [intensional approach], one asks the subject what he means by the given word or how he uses it; then one proceeds in Socratic fashion to test this first answer by confronting the subject with counterexamples and borderline cases, and so on until the subject settles down more or less permanently upon a definition or account (Mates, 1958, p. 165–166).

[6]Note, however, that Ryle and Austin are not, strictly speaking, disagreeing in the passages Mates cites because Ryle and Austin are discussing different words: the adjectives "voluntary" and "involuntary" in Ryle's case, and the adverbs "voluntarily" and "involuntarily" in Austin's case.

According to Mates, both the extensional and intensional approaches are equally relevant to verifying whether someone uses a word in a given way or with a given sense, and there is no guarantee that the two approaches yield the same results.[7] It is possible, for example, to observe a speaker using a word in a way that she says she would never use it.[8]

Mates says that ordinary language philosophers "tend toward an armchair version of the extensional method" of verifying statements about language, and suggests that they neglect the intensional method. Even if philosophers apply "voluntary" and "involuntary" in different ways than ordinary folk, the two groups may give the same "intensional" account of the word's meaning. So there may be evidence that pulls in the direction of thinking that philosophers use "voluntary" and "involuntary" differently than ordinary folk, and there may be equally significant evidence that pulls in the opposite direction:

> Thus, even if Professor Ryle had determined that ordinary folk in fact apply the word 'voluntary' only to actions which ought not to be done, while philosophers apply it to meritorious actions as well, he would be far from establishing that philosophers and ordinary folk apply the word 'voluntary' in different ways, i.e., attach different senses to it (Mates, 1958, p. 166).

Mates's challenge to ordinary language philosophy can be summarized as follows:

1. Statements about language made by ordinary language philosophers are descriptive (evaluable as true or false), not normative (i.e., not advice on how to use language).

2. Skepticism about the ordinary language philosophers' descriptive claims about language is warranted by the fact that (i) speakers are often not reliable reporters of their own linguistic behavior, and (ii) there is disagreement even among practitioners of ordinary language philosophy (Ryle and Austin, for example) about how expressions are used.

[7] See Hampton and Passanisi (2016) for a defense of the divergence of intensions and extensions in a contemporary theory of concepts.

[8] See Labov (1975) and Labov (2006) for examples of mismatches between speakers' judgments about how they speak and how they do in fact speak. For example, Labov discusses Philadelphia speakers who say they never use "anymore" in positive sentences, to mean roughly "nowadays", who then go on to use it that way (e.g., "Do you know what's a lousy show anymore? Johnny Carson") (Labov 1975, p. 107). And New Yorkers who were asked whether they pronounced "cards" with [r] or without [r] (*cahds*) overestimated how much of their own pronunciation was r-ful (Labov 2006, discussed in Gordon 2013, p. 69, from which the following passage is drawn):

> Labov recounts an especially poignant case of this phenomenon. He interviewed a middle class mother and her daughter who offered strong opinions about many of the phonological variables and ridiculed speakers from the subjective reaction test, including one who dropped an [r] in one word. They insisted that their own speech was impeccably r-ful. Labov "unwisely" played back part of their interview to demonstrate that they regularly dropped [r] too. This left them "disheartened in a way that was painful to see" (Labov 2006, p. 314)

.

3. There are two basic approaches to verifying descriptive statements about language: the extensional approach, which looks at the use of expressions, and the intensional approach, which involves eliciting the beliefs of speakers about the meaning or use of expressions. Ordinary language philosophers "tend toward an armchair version of the extensional approach", ignoring the intensional approach. The two approaches may yield conflicting results, so relying only on one method will not give a complete picture of ordinary use.

Before diving deep into Cavell's defense of the methods of ordinary language philosophy, I want to explore an alternative response to Mates's challenge that there are good reasons to reject. The failure of that defense will then show what Cavell's defense has to avoid in order to be successful.

3 How not to defend ordinary language philosophy: conventionalism and "legitimate use"

Both Searle (1969) and Vendler (1967) attempt to defend the idea of *a priori* knowledge of language against a Mates-style challenge. Both attempt to address point 2 in my reconstruction of Mates's challenge, by giving reasons to reject skepticism about the reliability of armchair claims about how expressions are used. Both Searle and Vendler respond to that kind of skepticism by drawing analogies between our knowledge of how to play games (baseball in Searle's case and chess in Vendler's) and our knowledge of how to use linguistic expressions.

Vendler explains how certain kinds of knowledge about chess don't involve conducting any empirical investigations:

> Suppose that while watching a game of chess I see two Pawns of the same color standing in the same column. Then I say: "One of them must have taken an opposing piece in a previous move". How do I know this? Is it sufficient to say that in all chess games we ever witnessed this correlation held? No, *given the rules of the game*, the relation holds *a priori*; the contrary is not something unusual or unlikely: it is inconceivable (Vendler, 1967, p. 17).

The linguist, according to Vendler, is analogous to an observer of chess who wants to describe its rules by observing what players of chess regard as admissible moves. Discovering the existence of the rules is an empirical project, but the rules lay down necessary requirements on movements in the game, and given knowledge of the rules, one can derive *a priori* knowledge about the possibility of particular configurations of pieces on the board. For Vendler, speakers of a language are entitled to statements about ordinary language on the same grounds that one who knows the rules of chess is entitled to the claim, made of two pawns of the same color standing in the same column that "one of them must have taken an opposing piece in a previous move". The entitlement is not based on observation, but derives from the speaker's knowledge of the rules of the language, which can be combined to derive knowledge of unobvious conclusions.

Similarly, Searle (1969) writes that his knowledge of particular facts about the language he speaks is analogous to his knowledge that after hitting a fair ball in baseball, "the batter runs in the direction of first base, and not in the direction, say, of third base or the left field grand stand" (p. 14). Searle asks, "Now what sort of knowledge is this? On what is it based? How is it possible?" He answers:

> My knowledge is based on knowing how to play baseball, which is *inter alia* having internalized a set of rules. I wish to suggest that my knowledge of linguistic characterizations is of a similar kind.
>
> If this is correct, then the answer to the philosopher's question, "What would we say if. . . ?" is not a prediction about future verbal behavior but a hypothetical statement of intention within a system of rules, where mastery of the rules dictates the answer . . .

The "game-like" account of how one knows statements about one's native language without having to observe any linguistic behavior goes as follows:

1. A speaker of a language knows a set of rules governing the language.
2. One can derive particular pieces of *a priori* knowledge about the language by applying those rules.

The "game-like" explanation of *a priori* linguistic knowledge proposed by Searle and Vendler is not radical; in fact it is a way of stating one of the basic commitments of standard compositional theories of semantics. Speakers of a language are able to derive the meaning of an infinite number of novel sentences from their knowledge of a finite stock of lexical items and rules for assembling those items into complex wholes. Such an account is meant to explain how speakers know the meaning of sentences they have never observed in use before; in that sense the standard semantic picture is a story about a form of *a priori* linguistic knowledge.

While this "game-like" explanation of how speakers can acquire knowledge of unobvious linguistic facts about the language they speak without leaving their armchairs looks promising, Jackman (2001) argues that such an account doesn't do the work the ordinary language philosopher needs it to do (Jackman calls this account "conventionalist", and I will adopt his terminology). The first prong of Jackman's attack involves driving a wedge between a speaker's *a priori* knowledge of her own idiolect and knowledge that *anyone else* speaks the same language. Jackman cites Searle saying "That my idiolect matches a given dialect group is indeed an empirical hypothesis (for which I have a lifetime of 'evidence'), but the truth that in my idiolect "oculist" means eye doctor is not refuted by evidence concerning the behavior of others" (Searle 1969, p. 13). The conventionalist story of how *a priori* linguistic knowledge (or knowledge of the rules of a game) is possible explains only how a speaker can have *a priori* knowledge of her own idiolect (or the rules of the game that she is following), but any move beyond the individual to knowledge of linguistic facts about a community of speakers (how "we" ordinarily use an expression) involves an "empirical hypothesis". The conventionalist therefore falls short of justifying statements of the sort that interest the ordinary language philosopher, which concern what *we* mean when we say certain things.

The second prong of Jackman's attack on the conventionalist account of *a priori* linguistic knowledge targets the assumption that speakers have *a priori* knowledge even of their own idiolect. Drawing on examples from Putnam (1975) and Burge (1979), Jackman argues that "a speaker's discriminatory capacities often underdetermine or even misidentify what one is talking about" (p. 318). For example, Burge describes a situation in which a competent speaker consistently *misapplies* the expression "arthritis", which reveals the possibility that speakers may have internalized mistaken linguistic rules. Recognizing the possibility of linguistic ignorance of the Putnam and Burge varieties "comes with a corresponding uncertainty about the mastery of one's own language" (p. 319).

The two prongs of Jackman's attack amount to a dilemma for someone who wants to defend the idea that we have *a priori* entitlement to claims about what "we" mean:

1. If speakers have *a priori* knowledge of what the expressions of their language mean, then they aren't entitled to generalize from their own language to a public language.

2. If the language that speakers know is shared, then they cannot have *a priori* knowledge of the meaning of expressions of the language.

The upshot of the dilemma is that there cannot be *a priori* knowledge that one's language (or the rules one is playing by) is shared—but that is just what it seems ordinary language philosophers need in order to respond to Mates's challenge.

Some recent defenders of the methods of ordinary language philosophy rely on versions of the conventionalist response to Mates that sidestep worries about idiolects by assuming that certain practitioners of ordinary language philosophy have access to "objective facts about linguistic norms". For example, Sandis (2010) distinguishes ordinary language philosophy from recent experimental investigations of topics of philosophical interest by claiming that ordinary language philosophy makes observations about "proper linguistic usage" and "*legitimate* use":

> OLP [ordinary language philosophy], by contrast [with experimental philosophy of language], is not interested in what the majority of people happen to think at any given time and place but, rather, in objective facts about linguistic norms (p. 185).
>
> When OLP talks of what we *would* say, it refers not to the latest up-to-the-minute statistics about what happens to be the case but to well-established, *legitimate* use, distinguishing further between paradigm and peripheral cases (p. 186).

The invocation of "correct" or "legitimate" use (or simply "the use"—see Hanfling (2000)) recalls Ryle's distinction between use and usage, where use is explicitly normative, while usage is not. When he considers this aspect of Ryle's view, Mates dismisses it, and says, "What authority deems it wrong to use the word 'voluntary' as the philosopher does [as opposed to the ordinary speaker]?" (Mates, 1958, p. 164) That is a centrally important

question for the defender of ordinary language philosophy who invokes "correct" or "legitimate" use. It is a question that Cavell attempts to answer, and his answer is strikingly different than the conventionalist response to Mates's challenge. I think V.C. Chappell (1964) is still correct (more than 50 years later) in his assessment that Cavell gives "the most detailed explanation and defense of the procedures of ordinary language philosophers that has yet appeared". In the remainder of this paper, I will describe Cavell's response to Mates's attack and show how it avoids the problems that arise for the conventionalist defense of ordinary language philosophy.

4 Cavell's defense of the methods of ordinary language philosophy in "Must We Mean What We Say?"

As Bates and Cohen (1972) observe, explicating Cavell "is no picnic", but the basic outlines of what Cavell argues in his (1958) defense of ordinary language philosophy are clear: He denies point 1 in Mates's challenge, by arguing that there is a normative element to statements about ordinary use, he disputes point 2 by arguing that native speakers of a language are authoritative regarding ordinary uses of expressions in that language, and he claims that Mates's discussion of intensional and extensional methods are "irrelevant" to the concerns of ordinary language philosophy as he defends it (Cavell, 1958, p. 131).

4.1 On the reliability of judgments about ordinary language

Cavell distinguishes two types of statements about ordinary language:

> Type 1: "Statements which produce *instances* of what is said in a language ('We do say . . . but we don't say —')"
>
> Type 2: Instances of what is said in a language "accompanied by *explications*—statements which make explicit what is implied when we say what statements of the first type instance us as saying ('When we say . . . we imply (suggest, say) —')".

The clash between Austin and Ryle involves statements of different types: Cavell says Austin produces a statement of type 1 ("we may make a gift voluntarily") while Ryle makes a claim of type 2, about what "voluntary" and "involuntary" mean—that when we describe an action as voluntary or involuntary, we imply that the action ought not to be done.[9] Cavell allows that Ryle's type 2 statement is wrong, but he rejects Mates's charge that Ryle lacks good evidence for making such a statement (whether or not it is correct). Ryle has good evidence for his claim simply in virtue of being a native speaker of English:

> Such speakers do not, in *general*, need evidence for what is said in the language; they are the source of such evidence. . .

[9]Cavell misdescribes Austin's remark, which is not metalinguistic—it doesn't mention "what we say", it simply produces an example of something that does make sense. The fact that it makes sense is evidence against Ryle's Type-2 statement, however, in the way Cavell describes.

> [F]or a native speaker to say what, in ordinary circumstances, is said when, no such special information is needed or claimed. All that is needed is the truth of the proposition that a natural language is what native speakers of that language speak (pp. 174–176).

Cavell is right that it would be absurd to think that native speakers of a language need "special information" or have to "count noses" before they know how to use some expression in the language they speak. There have to be native speakers saying things for any data about ordinary use to exist. But it's a different question when speakers go beyond speaking their language to make meta-linguistic claims about how expressions are ordinarily used. Fodor and Katz (1963, p. 60) criticize Cavell for missing the distinction between using an expression and making statements about how that expression is used:

> What Cavell misses is the distinction between what a native speaker says (the utterances he produces in the course of speaking) and what he says about what he and other native speakers say (the metalinguistic comments he makes when the reflective mood is upon him). There can be no doubt but that most (though definitely not all) of the utterances of a native speaker are utterances of the speaker's language. This truth is guaranteed by the truism that a natural language is what a native speaker of that language speaks. However, the statements that a native speaker makes about his language, his metalinguistic claims, need not be true in order for the linguist to have noses to count.

Fodor and Katz overextend their criticism of Cavell here—though a restricted version of their objection is important. It is standard practice in linguistics and philosophy of language to make use of meta-linguistic judgments of speakers as evidence for and against theoretical hypotheses.[10] In syntax, theorists make use of speakers' meta-linguistic judgments about the acceptability of various constructions, like the following:[11]

Please indicate which of the following two sentences sounds more acceptable:

(1) Ginny remembered to have brought the beer.

(2) Ginny remembered to bring the beer. (Sprouse et al., 2013, p. 237)

And in semantics and pragmatics, speakers are often asked to make truth value judgments about claims made in various contexts, or judgments about entailment relations. For example, as part of the evaluation of the semantics of "donkey sentences", Geurts (2002 p. 135–136) asked participants to judge whether the statement "Every boy that stands next to a girl holds her hand" is a true description of the situation depicted in Figure 1.

[10] This is pointed out, in response to Fodor and Katz, by Bates and Cohen (1972).

[11] "The data syntacticians predominately appeal to are speaker/hearers' intuitive judgements of acceptability and interpretability (suffice it to say, any relevant data are admissible; intuitions are simply readily available, not obviously misleading, and have no serious competitor as a data source)" (Collins, 2007, p. 887). For a summary of discussions of the role of meta-linguistic intuitions in syntax, see Sprouse (2013).

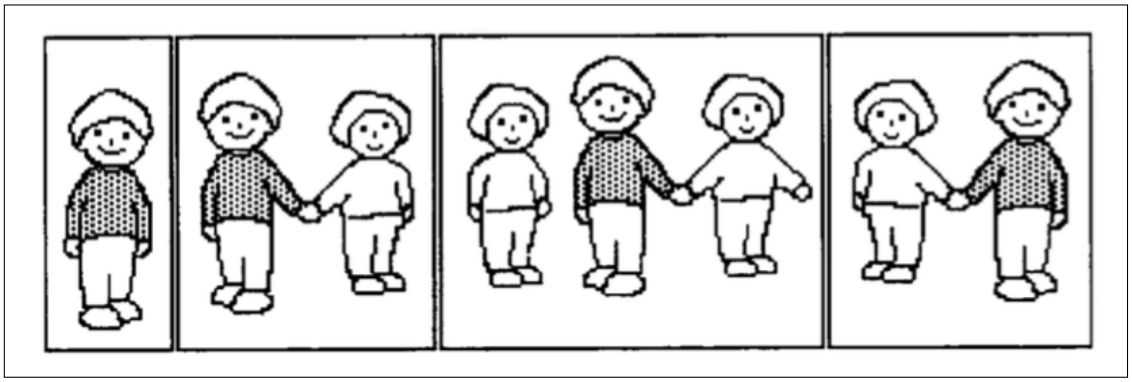

Figure 1: Every boy that stands next to a girl holds her hand.

The use of speakers' meta-linguistic judgments about truth and acceptability reflects a central assumption of contemporary research in the language sciences, namely, that such judgments are products of speakers' underlying (tacit) linguistic knowledge. While speakers' meta-linguistic judgments about acceptability and truth are widely regarded as useful tools in the investigation of syntax, meaning, and use, it is a further, much more contentious question whether individual speakers' judgments about how a *group* of speakers would use a particular expression are reliable.

Mates is skeptical that ordinary speakers are reliable when asked to make these kinds of judgments. Only recently have there been any systematic examinations that bear on whether ordinary speakers are reliable judges about other speakers' linguistic behavior. The studies that have been conducted have yielded results that, while they do not support general skepticism about the reliability of group-directed meta-linguistic judgments, also do not support Cavell's claim that the ordinary language philosopher has "good evidence" for statements of type 1 or type 2 simply in virtue of being a native speaker of English.

Experimental investigations in both linguistics and philosophy have indicated that meta-linguistic judgments about *acceptability* made by ordinary speakers and those made by experts are highly correlated with one another, which is evidence that generalizations from one's own judgment to judgments of the group are reliable (Culbertson and Gross, 2009; Myers, 2009; Sprouse et al., 2013). These findings might be taken to support the ordinary language philosophers' method of relying on armchair meta-linguistic judgments, since (in the acceptability judgment tasks that have been experimentally examined) those armchair judgments correlate with ordinary speakers' judgments in experimental conditions. But it is a further question whether ordinary and expert meta-linguistic judgments about *semantic* and *pragmatic* facts, like what particular names or natural kind terms refer to, correlate in the same way. And there is evidence from experimental investigations that, in certain cases, there is substantial interpersonal variation in meta-linguistic judgments.[12]

[12]For evidence concerning proper names, see the debate surrounding Machery et al. (2004). Genone (2012) and Hansen (2015) provide summaries of that debate. For evidence concerning natural kind terms, see Braisby et al. (1996), Jylkkä et al. (2009) Genone and Lombrozo (2012), and Nichols et al. (2016). Hansen (2015) provides a summary of the experimental evidence for natural kind terms. For related evidence

In summary, given the existing experimental evidence, there is reason to think that armchair meta-linguistic judgments about certain linguistic properties (acceptability) reliably generalize, while armchair meta-linguistic judgments about other properties (reference, satisfaction conditions for adjectives) may not. Overall, that is bad news for both general skepticism about the reliability of armchair metalinguistic judgments *and* for Cavell's optimism about the reliability of such judgments. It will take much more experimental work to determine the exact contours of ordinary speakers' ability to make reliable judgments about the linguistic behavior of groups they belong to.

4.2 Categorial declaratives

Cavell's response to Mates's challenge that invokes the fact that native speakers of a language are reliable judges of how words are ordinarily used does not stand up to scrutiny. But there is a different line of response suggested by Cavell that isn't subject to the same worries. I will develop that response in the remainder of this paper.

Cavell (1958) is centrally concerned with the relation between what someone *says* and what she *means* by what she says. The significance of the question in the title of the essay, "Must We Mean What We Say?", is whether there is any logical (or otherwise necessary) connection between "what you (explicitly) say and what saying it implies or suggests" (p. 196). Cavell revises Ryle's type 2 statement about "voluntary" to accommodate Austin's observations, and produces the following statement, which he refers to as "S":

S "When we ask whether an action is voluntary we imply that the action is fishy"

Cavell says that "Something important will be learned if we realize that we do *not* know what kind of assertion S is" (p. 181). S is "obviously" not analytic, Cavell says, because "it is perfectly true that 'voluntary' does not *mean* (you will not find set beside it in a dictionary) 'fishy')" (p. 181). And yet Cavell expresses the feeling that S expresses a necessary truth, and he assumes that if a statement expresses a necessary truth then it has to be known *a priori*. That combination of necessity and *a prioricity* with non-analyticity is what poses the problem Cavell tries to solve:

> When (if) you feel that S is necessarily true, that it is *a priori*, you will have to explain how a statement which is obviously not analytic can be true *a priori* (p. 181).[13]

It is important to note that Cavell is writing at a time before widespread acknowledgement of two conceptual distinctions that became very important to analytic philosophy in the latter half of the 20th century:

- Necessary truths can be known *a posteriori*.

concerning adjectives, see Hansen and Chemla (2016).

[13]Cavell was teaching the *Critique of Pure Reason* at Berkeley around the time that he was writing "Must We Mean What We Say?" (Cavell, 2010, p. 277).

- There is a distinction between what a statement necessarily *implies* in virtue of its (context-invariant) meaning and what making a statement *implicates*, which is context dependent.

Acknowledging either distinction would help explain what kind of assertion S is. First, supposing that S does express a necessary truth, it need not be knowable *a priori*. Informative identity statements, like "Hesperus is Phosphorus" are necessarily true, but can be *a posteriori* discoveries (Kripke, 1980, Lecture II). Even if Cavell were right that S expresses a necessary truth, it could be discovered via an empirical scientific investigation of language.

Second, Cavell is interested in the regularities that exist between what is explicitly said and what is suggested or conveyed by what is said. It is these regularities that are of central importance to the ordinary language philosophers who want to argue that philosophers are "stretching" or distorting the ordinary meaning of expressions by ignoring the connections between what is explicitly said and what is conveyed by what is said. Ryle's discussion of the way philosophers distort the meaning of "voluntary" and "involuntary", for example, relies on the idea that there is some regular connection between those expressions and the actions they describe being something that someone ought not to do. That is, Ryle thinks that saying that some action is voluntary will convey that the action ought not to be done. When philosophers wonder whether actions of a certain type are voluntary or not when there is no question of whether or not the action ought not to be done, then they are distorting the ordinary meaning of "voluntary".

Cavell calls the connection between what is explicitly said and what is suggested or conveyed by what is said a "necessary" connection.[14] But that is to overstate the kind of regularity that exists between what is explicitly said and what is conveyed or suggested by what is explicitly said. Consider Cavell's own example of a "necessary" truth, namely S. It is doubtful that S expresses a necessary truth, because there are contexts in which one can ask whether an action is "voluntary" without suggesting that there is anything fishy about the action (bracketing the fact that it's not really clear what a "fishy" action would be). Consider the following questions, which are based on naturally occurring sentences obtained through the Corpus of Contemporary American English (COCA):

(3) Was the contraception proposed by the Gates Foundation as part of the Family Planning 2020 initiative going to be voluntary? (`http://goo.gl/gkzpHH`)

(4) Was the administrative fee taken out during the first three years of employment voluntary? (`http://goo.gl/axQqCp`)

It is unclear what the "fishy" suggestion either of these questions would be. The first question concerns whether the contraception is voluntary, as opposed to mandatory, and the second question concerns whether the administrative fee can be waived by the employee or not.

A better way of capturing the relation between what is said and what is conveyed by what is said is introduced by Grice's (1975) "logic of conversation", which wouldn't have

[14] See Mulhall (1999, pp. 13–14) for further discussion.

been available to Cavell at the time of writing "Must We Mean What We Say?".[15] Conversational implicatures arise out of a combination of the content of what is explicitly said and the "cooperative principle"—a governing principle of conversation. For example, if I ask whether some action is voluntary, I implicate (though I do not explicitly say) that I don't know whether the action is voluntary. Assuming that I am making a cooperative contribution to the conversation, my audience can (with reason) infer that I don't know whether the action was voluntary or not. That they draw that inference is something that I can expect if I understand the norms of conversation. I can exploit this expectation by asking whether some action was voluntary when I know that it was voluntary; that will generate the false suggestion that I don't know whether the action was voluntary. It is the norms governing cooperative conversations, rather than anything about the meaning of "voluntary", that generate expectations regarding the question. Expectations about what I have implicated are not completely under the speaker's control, but they also do not reflect necessary connections between what is explicitly said and what is conveyed or suggested. To the contrary: implicatures differ from the context-invariant implications associated with what is explicitly said in that they are context-dependent.

The unavailability of those two conceptual distinctions constrains the range of options open to Cavell in his discussion of statements about ordinary language, and he recognizes that the conceptual tools available to him are inadequate to his task:

> At this point the argument has become aporetic. 'Statements about ordinary language'... are not analytic, *and* they are not (it would be misleading to call them) synthetic (just like that). Nor do we know whether to say they are *a priori*, or whether to account for their air of necessity as a dialectical illusion, due more to the motion of our argument than to their own nature. Given our current alternatives, there is no way to classify such statements; we do not yet know what they are (p. 184).

Even given these limitations, part of Cavell's proposal for understanding "statements about ordinary language" is worth salvaging. In the remaining part of this section I will spell out what I think that part is.

Cavell begins his discussion of how it is possible to be entitled to make S-type statements without conducting empirical observations by distinguishing S from a similar-sounding statement T, which he says "does seem obviously synthetic":

T "Is X voluntary?" implies X is fishy.

He says:

> But S and T, though they are true together and false together, are not everywhere interchangeable; the identical state of affairs is described by both, but a

[15] The lectures that were the basis of Grice (1975), "Logic and Conversation", were delivered in 1967. See Grice (1975, p. 41). It is very doubtful that Cavell himself would endorse this way of characterizing the relation.

> person who may be entitled to say T, may not be entitled to say S. Only a native speaker of English is entitled to the statement S, where as a linguist describing English may, though he is not a native speaker of English, be entitled to T. What entitles him to T is his having gathered a certain amount and kind of evidence in its favor. But the person entitled to S is not entitled to *that* statement for the same reason. He *needs* no evidence for it... The question of evidence is irrelevant (p. 182).

The interesting suggestion in this passage is that one can be entitled to make a statement in virtue of having the right kind of status as a member of a certain group—in this case, the group of native speakers of English. S, unlike T, uses the first person plural "we", which presupposes that the speaker is included in the group (or "plural individual") denoted by "we" (Schlenker, 2003, p. 5). If the presupposition is satisfied, it can be the case that the speaker is entitled to "speak for the group", and introduce, or reinforce, a *rule*.

Cavell observes that statements that share the same surface form as type-T (or type-S) statements can be used to perform a different kind of speech act[16]—instead of purporting to state a fact, they can be used to introduce a rule, "a way of doing or saying something which is to be *followed*":

> Whether remarks... 'about' ordinary language, and equally about ordinary actions—are statements or rules depends on how they are taken: if they are taken to state facts and are supposed to be believed, they are statements; if they are taken as guides and supposed to be followed, they are rules (p. 184).

Statements about what "we" do can play this role: "[I]f someone is tempted not to fulfill a promise, you may say, 'Promises are kept', or 'We keep our promises (that is the sort of thing a promise is)', thus employing a rule-description—what I have called a categorial declarative" (p. 195). A categorial declarative would not be challenged by citing evidence that the group referred to by "we" does not always conform to the proposed rule. Consider the following sentence:

(5) We send thank-you cards in this family.

The utterer of (5), understood as a categorial declarative, is not entitled to it on the grounds of observing the past behavior of her family. It may even turn out that members of the family have not usually sent thank-you cards. But that fact does not mean that the speaker is not entitled to utter (5), and pointing it out would not constitute a refutation of (5). Understood as a categorial declarative, (5) could be challenged in the following way:

(6) That's not for you to say.

The entitlement to the utterance (5) in the situation described does not stem from observation, but from occupying a certain position of authority in the group spoken for. As

[16]Bates and Cohen (1972, p. 11) observe that "Cavell is concerned with something like speech acts [in his discussion of categorial declaratives]".

Cavell says about statements of type S, understood as categorial declaratives, the question of evidence is irrelevant for one's entitlement to assert (5).

The parallel between (5) and the case of claims about what "we" say is that speakers can occupy an analogous position of authority among fellow speakers of a language. Statements like S can be used to perform two different kinds of speech act. Critics of ordinary language philosophy (Mates and Fodor and Katz, for example) have focused exclusively on the act of describing, and wonder where the entitlement for such a claim comes from if not from observation of actual linguistic behavior. But Cavell's central claim in "Must We Mean What We Say?", as I understand it, is that the rule-giving aspect of statements like S has been overlooked. Entitlement to the rule-giving kind of speech act is derived from occupying a position of suitable authority with respect to those included in the group referred to with "we", and challenging that kind of speech act requires a different kind of challenge than presenting evidence that "we" do not in fact act as the speaker seems to be saying that we do. Challenging the speaker's entitlement to the rule-giving aspect of the categorial declarative involves contesting the speaker's entitlement to speak on behalf of the relevant group (as in (6)).

If this way of understanding Cavell's proposal is correct, it would provide a way of responding to Mates's challenge that doesn't run into objections from the spotty reliability of armchair meta-linguistic judgments about how we speak. Recall that Mates rejects the idea that the statements of ordinary language philosophers about language are "normative", even though Ryle distinguishes what he calls "use" and "misuse" from the descriptive patterns of "usage". Mates's criticisms are directed at the statements of ordinary language philosophers on the assumption that they are (non-normative) descriptions. But if my interpretation of Cavell is correct, then Cavell should be seen as rejecting Mates's starting assumption and insisting that there is a central normative component to the statements of the ordinary language philosophers—statements about what we say are not just descriptions, they are proposals as to how words should be used.[17]

4.3 Aesthetic judgments and claims about ordinary language

In 'Must We Mean What We Say?', Cavell does not discuss in detail what would entitle someone to make S-type statements about ordinary language. He says "An examination of what does entitle a person to the statement S would be required in any full account of such statements. Such an examination is out of the question here" (Cavell 1958, p. 182). But he returns to the topic and enriches his discussion with a new metaphor in 'Aesthetic Problems of Modern Philosophy' (Cavell 1976a, originally published in 1965). His new

[17] Some early commentators on ordinary language philosophy detected the normative component in statements about ordinary language, but without developing the idea in any detail: "We strongly suspect that cases of putative ordinary language analysis are, in fact, disguised reformations" (Maxwell and Feigl, 1961, p. 489); "There are innumerable cases in the philosophy of language where we realize that what we off hand may have interpreted to be a language hypothesis is... more readily understood as a verbal recommendation, as a convention, or as any other type of normative statement, say a proposal for how to use a given linguistic expression..." (Tennessen, 1962, p. 507).

way of thinking about 'the characteristic claim. . . of ordinary language philosophers' is to compare them with aesthetic judgments:

> I will suggest that the aesthetic judgment models the sort of claim entered by these philosophers, and that the familiar lack of conclusiveness in aesthetic argument, rather than showing up an irrationality, shows the kind of rationality it has, and needs (Cavell, 1976a, p. 86).

Cavell uses different patterns of disagreement and reason-giving to illustrate the difference between two types of aesthetic judgment—a mere judgment of "sense" versus a judgment of "reflection"—and a non-aesthetic judgment (whether there is a goldfinch in the garden). The patterns of disagreement that distinguish the two types of aesthetic judgment go as follows:

1. A: Canary wine is pleasant.
   B: How can you say that? It tastes like canary droppings.
   A: Well, I like it.
2. A: He plays beautifully, doesn't he?
   B: Yes, too beautifully. Beethoven is not Chopin.

Or he [B] may answer:

> B: How can you say that? There was no line, no structure, no idea what the music was about. He's simply an impressive colorist (p. 91).

Whereas it is unremarkable to defend a mere judgment of sense by responding to a challenge with "Well, I like it" (as in dialogue 1), it would be "feeble" to respond to the challenge in dialogue 2 by saying "Well, I liked it". Moreover, while A doesn't have to engage with B by giving some reasons to defend the assessment of the beautiful playing, if he doesn't, Cavell says "there is a price he will have to pay in our estimate of him" (p. 92). That is a difference in the obligations that one undertakes when making a judgment of reflection rather than a mere judgment of sense. Cavell then uses the same strategy to distinguish both types of aesthetic judgment from a judgment about a (more) straightforward matter of fact—whether there is a goldfinch in the garden:

A: There is a goldfinch in the garden.
B: How do you know?
A: From the color of its head.
B: But goldcrests also have heads that color.
A: Well, *I* think it's a goldfinch (it's a goldfinch to me).

In the goldfinch discussion, it's not open to A to retreat to a personal judgment without giving up his claim to be a "competent interlocutor in matters of knowledge (about birds?)" (p. 92).

In addition to the different patterns of challenge and retraction that characterize judgments of reflection, Cavell ties the success of such judgments not to their accuracy in stating the facts, but to their ability to convince the audience to see and acknowledge the relevant fact. The story of Sancho Panza's relatives Hume relates in "Of the Standard of Taste" is used to illustrate this criterion of success: they are ridiculed for their evaluation of the quality of a hogshead of wine on the grounds that they say it tastes of iron and leather.[18] But when the barrel is empty, a rusty iron key on a leather thong is found at the bottom. Cavell's reading of this episode is that it is an example of the *wrong* model of understanding how aesthetic judgments are vindicated:

> It dissociates the exercise of taste from the discipline of accounting for it: but *all* that makes the critic's expression of taste worth more than another man's is his ability to produce for himself the thong and key of his response; and his vindication comes not from his pointing out that it is, or was, in the barrel, but in getting us to taste it there (p. 87).

How should this model be applied to S-type statements about ordinary language? It gives a fuller picture of the thin nature of authority that is required for someone to be entitled to S-type statements about ordinary language that remains underdescribed in "Must We Mean What We Say?" While the authority that is required to *make* such statements consists in merely being a speaker of the relevant language, the success conditions of such statements involve a very demanding, and fickle, component: successful categorial declaratives require getting one's audience to acknowledge their correctness by adopting the proposed rule of use. Those elements clearly distinguish categorial declaratives from statements about ordinary use made by linguists: entitlement to such statements does require possessing evidence of how people do in fact use language, and the success of such statements does not require the audience to embrace any rule of use.

4.4 Austin at criticism

For Cavell, Austin represents the "purest version of 'ordinary language philosophy'" (Cavell, 1976b, p. 98), and he compares Austin's practice of drawing fine distinctions and making illuminating comparisons to the methods of an art critic:

> The positive purpose in Austin's distinctions resembles the art critic's purpose in comparing and distinguishing works of art, namely, that in this crosslight the capacities and salience of an individual object are brought to attention and focus (Cavell, 1976b, p. 103).

Examining how Austin distinguishes and compares the meaning of expressions will make Cavell's comparison of Austin to an art critic—and more generally, his comparison of statements about ordinary language to aesthetic judgments—more concrete.

[18] Presumably those were not accepted terms for describing the taste of wine in 17th century Spain; maybe a more contemporary equivalent would be saying that a wine tastes of "too many tramlines". See the "Winesmanship" section of Potter (2005) for discussion.

Austin's most famous lexical distinction concerns the expressions "mistake" and "accident". He says that the choice to apply one or the other of these expressions to describe a situation can "appear indifferent... Yet a story or two, and everybody will not merely agree that they are completely different, but even discover for himself what the difference is and what each means" (Austin, 1957, pp. 10–11). To distinguish the meaning of these expressions, Austin tells two contrasting "donkey stories", which make it seem clear that "by mistake" better describes the action in one situation, and "by accident" better describes the action in the other, thereby providing evidence that the meanings of the two expressions are indeed distinct:

> You have a donkey, so have I, and they graze in the same field. The day comes when I conceive a dislike for mine. I go to shoot it, draw a bead on it, fire: the brute falls in its tracks. I inspect the victim, and find to my horror that it is your donkey. I appear on your doorstep with the remains and say—what? 'I say, old sport, I'm awfully sorry, &c., I've shot your donkey by accident?' Or 'by mistake'? Then again, I go to shoot my donkey as before, draw a bead on it, fire—but as I do so, the beasts move, and to my horror yours falls. Again the scene on the doorstep—what do I say? 'By mistake'? Or 'by accident'? (Austin 1957, p. 11 n. 4)

Those who read Austin's donkey stories tend to reach the same conclusion about which term to apply in which situation: "by mistake" better suits the first story, and "by accident" the second.[19] Backed by the donkey stories, we can issue two S-type statements that have achieved the kind of vindication, in their acknowledgement by those who read Austin's donkey stories, that Cavell says is an aim of aesthetic judgments:

(7) In the first story, we say you shot your neighbor's donkey by mistake.

(8) In the second story, we say you shot your neighbor's donkey by accident.

A comparison with a less successful application of Austin's method shows, by way of contrast, how clearly the donkey stories compel us to acknowledge a difference in meaning between doing something "by mistake" and "by accident". The gocart story from (Austin, 1966) is intended to distinguish the meaning of "intentionally" and "deliberately":

> I am summoned to quell a riot in India. Speed is imperative. My mind runs on the action to be taken five miles down the road at he Residency. As I set off down the drive, my cookboy's child's new gocart, the apple of her eye, is right across the road. I realize I could stop, get out, and move it, but to hell with that: I must push on. It's too bad, that's all: I drive right over it and am on my way. In this case, a snap decision is taken on what is essentially an incidental matter. I did drive over the gocart deliberately, but not intentionally—nor, of course, unintentionally either. It was never part of my intention to drive over the gocart. At no time did I intend to drive over it. It was incidental to anything

[19] See, for example, Gustafsson (2005, p. 368), and Hanfling (2000, p. 64).

> I intended to do, which was simply to get to the scene of the riot in order to quell it. However "odd" it may sound, I feel little doubt that we should say here that we did run over the gocart deliberately and that we should not care to say we ran over it intentionally. We never intended to run over it (p. 432).

One crucial difference between this story and the donkey stories is that Austin provides a "gloss" on how to describe the action of driving over the gocart: he says that it was done deliberately, but not intentionally, because "we never intended to run over it". Prompted by the feeling that this story didn't compel agreement to the same degree as the donkey stories, and the fact that it has been argued that doing something intentionally does not require *intending* to do it (see Harman 1976), Hansen and Chemla (2015) constructed a formal experiment in which the gocart story was presented to experimental participants in three conditions: (i) with Austin's original gloss, (ii) with a reversed gloss that made the case that the gocart was driven over intentionally, but not deliberately, and (iii) with no gloss either way. In Hansen and Chemla's experiment, it was easy to reverse participants' judgments about the gocart story by reversing the glosses, and in the "no-gloss" condition, where participants were asked to say whether "deliberately" or "intentionally" better described the action of driving over the gocart, no preference was observed either way. (Fig. 1, Hansen and Chemla 2015). The experiment lends support to the thought that the story itself (without a gloss) is insufficient to generate agreement with Austin's judgement that we should say that we ran over the gocart deliberately, but not intentionally.

If the criteria of success for an S-type statement of "what we say" were the same as those for an observation made by an empirically-informed linguist, then Austin's remarks regarding how we use "deliberately" and "intentionally" would simply be incorrect. The gocart experiment in Hansen and Chemla (2015) did not find any evidence that English speakers have a preference regarding whether "deliberately" and "intentionally" better describes the action of running over the gocart. But if we follow Cavell in his analogy between S-type statements and aesthetic judgments, Austin's gloss is *not* a detachable component of the case for why we (should) describe the story in the way Austin suggests. In his reading of the gocart story, Austin makes the case that we should describe the action of running over the gocart as done deliberately, but not intentionally, because we did not *intend* to run it over. As mentioned above, whether that connection exists has been contested. But the important point for assessing Austin's methodology is that in his account of what "we should say" or "should not care to say" about the case, he is proposing that "deliberately" is a better term to use in describing the action than "intentionally"—whether or not speakers would prefer to use one term or the other before adopting his proposal.

4.5 Categorial declaratives and meta-linguistic negotiation

Cavell's analogy between categorial declaratives (S-type statements) and aesthetic judgments has provided a sketch of what is required for one to be entitled to make such statements (merely being a speaker of the language) and what is required for such statements to be successful (audience endorsement of the proposed use). Our understanding of how

categorial declaratives operate can be enriched by considering recent discussions in philosophy of language of how what look like straightforward statements of fact can be playing a normative, metalinguistic role (Sundell, 2010; Plunkett and Sundell, 2013, 2014).

Barker (2002, pp. 1–2) describes a "metalinguistic use" of "tall":

> Normally, (9) will be used in order to add to the common ground new information concerning Feynman's height:
>
> (9) Feynman is tall.
>
> But (9) has another mode of use. Imagine that we are at a party. Perhaps Feynman stands before us a short distance away, drinking punch and thinking about dancing; in any case, the exact degree to which Feynman is tall is common knowledge. You ask me what counts as tall in my country. "Well," I say, "around here,..." and I continue by uttering (9).

Barker (p. 2) characterizes what is communicated by an utterance of (1) in the situation he describes as follows: "My purpose in uttering (9) under such circumstances would be nothing more than to communicate something about how to use a certain word appropriately—it would be a metalinguistic use".

Plunkett and Sundell enrich Barker's picture of metalinguistic use by adding the possibility that metalinguistic uses of expressions can debate not only how expressions *are* used, but how those expressions *should be* used (Plunkett and Sundell, 2013, p. 3). They illustrate this possibility with metalinguistic disputes over the use of evaluative expressions like "spicy", where it isn't the case that there is some antecedently settled local standard for what makes something count as spicy. Spiciness is up for grabs in such a situation, and the appropriate use of "spicy" has to be "negotiated" by the participants in the conversation (p. 15). A pair of statements like (10) and (11) can function as moves in a metalinguistic negotiation about the appropriate use of the adjective "spicy":

(10) That chili is spicy!

(11) No, it's not spicy at all. (Plunkett and Sundell, 2013, p. 15)

Adopting Plunkett and Sundell's terminology, let a "metalinguistic proposal" be an individual move in a metalinguistic negotiation: it is a proposal for how an expression should be used.

A metalinguistic proposal, like a categorial declarative, has a descriptive form and a normative conversational role. Cavell's central examples of categorial declaratives involve statements about what saying something conveys or implicates, and what is implicated by what is said should be a good candidate for metalinguistic negotiation, since they aren't explicitly encoded in what is said and Grice observes that what is implicated is indeterminate

(Grice, 1975, p. 58).[20] There should therefore be room to negotiate what is conveyed or implicated by an utterance.

Consider again Cavell's example of a categorial declarative, S:

S "When we ask whether an action is voluntary, we imply that the action is fishy"

How can a speaker be entitled to say S without "counting noses"? Once S is understood as a metalinguistic proposal for how "voluntary" should be used (as implicating that the action it describes is somehow "fishy"), and not a description of the antecedently settled use of "voluntary", then the entitlement to S will come from the speaker's status as a speaker of the language. But the success of such a proposal will depend on the speaker's ability to get her audience to accept her proposal, as Cavell discusses in his invocation of the parallel between statements about ordinary use and aesthetic judgments.

Using the notion of a metalinguistic proposal, it is also possible to illuminate what Austin is doing when he draws fine-grained differences in meaning between expressions like "intentionally" and "deliberately". As discussed above, Hansen and Chemla (2015) found no evidence that ordinary speakers draw a clear distinction between the meaning of doing something "intentionally" and doing it "deliberately". But Austin wants to distinguish the meaning of those two expressions, and he argues that in his gocart story, it is better to say that the protagonist runs over the gocart deliberately, but not intentionally, because the protagonist didn't *intend* to run over the gocart. Austin, in discussing his method for distinguishing the meaning of "intentionally" and "deliberately" says:

> We must... imagine some cases (imagine them carefully and in detail and comprehensively) and try to reach agreement upon what we should in fact say concerning them. (Austin, 1966, p. 429)

Austin's methods for reaching agreement include "glossing" those cases—making metalinguistic proposals for how the relevant expressions should be used.

5 Must we measure what we mean?

Cavell argues that speakers of a language can be entitled to claims about how "we" use expressions, without having to "count noses". Mates (1958), Fodor and Katz (1963), and Jackman (2001) have criticized various attempts to explain the source of that entitlement. I have argued that an essential component of defending Cavell's proposal requires emphasizing the normative nature of "categorial declaratives". More specifically, I have suggested that categorial declaratives should be understood as proposals for how expressions should be used, or for what we should understand to be implicated by a statement. One can be entitled to such claims in virtue of being a speaker of the relevant language without having

[20] Grice's claim about the indeterminacy of implicatures was brought to my attention by Gray (2016). Plunkett and Sundell (2013, p. 3) say that disputes over how words should be used could arise when we are 'resolving ambiguity, prescisifying a vague term, setting a contextual parameter, or in any other way determining how some antecedently indeterminate matter of meaning should be settled'.

to conduct empirical observations of how ordinary people do in fact speak. But Cavell's proposal for how to understand the success conditions of such claims is both demanding and fragile—it is analogous to the conditions of a successful aesthetic judgment, which, according to Cavell, depends on the audience having the appropriate kind of reaction (analogous to getting the audience to *taste* the leather and iron in the wine in the story about Sancho Panza's relatives). That condition on the success of a categorial declarative is hard to satisfy—Cavell's own attempt to tie uses of "voluntary" to *fishiness* is not convincing, for example, for the reasons discussed in §4.2.

So the answer to the question *must we measure what we mean?*—when that is taken as a question about the need to survey how speakers of the language actually speak when making claims about what we mean by the use of an expression—is *no*, but it is extremely difficult to make a compelling proposal about the meaning of an expression that ignores evidence about how the expression is in fact used. Citing experimental (or corpus-based) evidence of a distinction in use can be an effective way of convincing an audience that a particular difference in meaning exists, and can play a powerful rhetorical role in producing the acknowledgement necessary for a successful categorial declarative. Ryle referred to this as the feeling of being "strengthened, when told that big battalions are on [your] side" (Ryle, 1953, p. 177). Even though Cavell is right that it is possible to be entitled to make claims about how "we" use certain expressions without engaging in experimental or corpus-based investigations of how words are used, contemporary practitioners of ordinary language philosophy should not ignore the power of the big battalions.

References

Austin, J.L. (1956–1957). "A Plea for Excuses". *Proceedings of the Aristotelian Society 57*, 1–30.

Austin, J.L. (1962). *Sense and Sensibilia*. Oxford: Oxford University Press.

Austin, J.L. (1966). "Three Ways of Spilling Ink". *The Philosophical Review 75*(4), 427–440.

Barker, Chris (2002). "The Dynamics of Vagueness". *Linguistics and Philosophy 25*(1), 1–36.

Bates, Stanley and T. Cohen (1972). "More on What We Say". *Metaphilosophy 3*(1), 1–24.

Baz, Avner (2012). *When Words Are Called For: A Defense of Ordinary Language Philosophy*. Cambridge, Massachusetts: Harvard University Press.

Braisby, Nick, B. Franks, and J. Hampton (1996). "Essentialism, Word Use, and Concepts". *Cognition 59*(3), 247–274.

Cavell, Stanley (1958). "Must We Mean What We Say?". *Inquiry 1*, 172–212.

Cavell, Stanley (1976a). "Aesthetic Problems of Modern Philosophy". In *Must We Mean What We Say?*, pp. 73–96. Cambridge: Cambridge University Press.

Cavell, Stanley (1976b). "Austin at Criticism". In *Must We Mean What We Say?*, pp. 97–114. Cambridge: Cambridge University Press.

Cavell, Stanley (2010). *Little Did I Know: Exercepts from Memory*. Stanford, CA: Stanford University Press.

Chapman, Siobhan (2011). "Arne Naess and Empirical Semantics". *Inquiry 54*(1), 18–30.

Chappell, V.C. (Ed.) (1964). *Ordinary Language*. Englewood Cliffs, New Jersey: Prentice-Hall.

Collins, John (2007). "Linguistic Competence without Knowledge of Language". *Philosophy Compass 2*(6), 880–895.

Culbertson, Jennifer and S. Gross (2009). "Are Linguists Better Subjects?". *British Journal of Philosophy of Science 60*(4), 721–736.

DeRose, Keith (2005). "The Ordinary Language Basis for Contextualism, and the New Invariantism". *The Philosophical Quarterly 55*(219), 172–198.

Fischer, Eugen (2014). "Verbal Fallacies and Philosophical Intuitions: The Continuing Relevance of Ordinary Language Analysis". In B. Garvey (Ed.), *J.L. Austin on Language*, pp. 124–140. Basingstoke: Palgrave MacMillan.

Fodor, Jerry A. and J. J. Katz (1963). "The Availability of What We Say". *The Philosophical Review 72*(1), 57–71.

Genone, James (2012). "Theories of Reference and Experimental Philosophy". *Philosophy Compass 7*(2), 152–163.

Genone, James and T. Lombrozo (2012). "Concept Possession, Experimental Semantics, and Hybrid Theories of Reference". *Philosophical Psychology 25*(5), 717–742.

Gordon, Matthew J. (2013). *Labov: A Guide for the Perplexed*. London: Bloomsbury.

Gray, Aidan (2016). "What's Wrong with Indeterminate Implicature?". Unpublished ms.

Grice, H.P. (1961). "The Causal Theory of Perception". *Proceedings of the Aristotelian Society 35*, 121–168.

Grice, H.P. (1975). "Logic and Conversation". In P. Cole and J. L. Morgan (Eds.), *Syntax and Semantics*, Volume 3, pp. 41–58. New York: Academic Press.

Gustafsson, Martin (2005). "Perfect Pitch and Austinian Examples: Cavell, McDowell, Wittgenstein, and the Philosophical Significance of Ordinary Language". *Inquiry 48*(4), 356–389.

Hampton, James A. and A. Passanisi (2016). "When Intensions Do Not Map Onto Extensions: Individual Differences in Conceptualization". *Journal of Experimental Psychology: Leaning, Memory, and Cognition 42*(4), 505–523.

Hanfling, Oswald (2000). *Philosophy and Ordinary Language: The Bent and Genius of our Tongue*. London: Routledge.

Hansen, Nat (2014). "Contemporary Ordinary Language Philosophy". *Philosophy Compass 9*(8), 556–569.

Hansen, Nat (2015). "Experimental Philosophy of Language". *Oxford Handbooks Online DOI: 10.1093/oxfordhb/9780199935314.013.53*.

Hansen, Nat and E. Chemla (2015). "Linguistic Experiments and Ordinary Language Philosophy". *Ratio 28*(4), 422–445.

Hansen, Nat and E. Chemla (2016). "Color Adjectives, Standards and Thresholds: An Experimental Investigation". Forthcoming in *Linguistics and Philosophy*.

Harman, Gilbert (1976). "Practical Reasoning". *The Review of Metaphysics 29*(3), 431–463.

Henson, Richard G. (1965). "What We Say". *American Philosophical Quarterly 2*(1), 52–62.

Jackman, Henry (2001). "Ordinary Language, Conventionalism and *a priori* Knowledge". *dialectica 55*(4), 315–325.

Jylkkä, Jussi, H. Railo, and J. Haukioja (2009). "Psychological Essentialism and Semantic Externalism: Evidence for Externalism in Lay Speakers' Language Use". *Philosophical Psychology 22*(1), 37–60.

Kripke, Saul (1980). *Naming and Necessity*. Cambridge, Massachusetts: Harvard University Press.

Labov, William (1975). "Empirical Foundations of Linguistic Theory". In R. Austerlitz (Ed.), *The Scope of American Linguistics: Papers of the First Golden Anniversary Symposium of the Linguistic Society of America*, pp. 77–134. Lisse: The Peter de Ridder Press.

Labov, William (2006). *The Social Stratification of English in New York City* (2nd ed.). Cambridge: Cambridge University Press.

Laugier, Sandra (2013). *Why We Need Ordinary Language Philosophy*. Chicago: University of Chicago Press.

Lyas, Colin (Ed.) (1971). *Philosophy and Linguistics*. Macmillan.

Lynch, Jack (2009). *The Lexicographer's Dilemma*. New York: Walker & Company.

Machery, Edouard, R. Mallon, S. Nichols, and S. P. Stich (2004). "Semantics, Cross-Cultural Style". *Cognition 92*, B1–B12.

Mates, Benson (1958). "On the Verification of Statements about Ordinary Language". *Inquiry 1*(1), 161–171.

Maxwell, Grover and H. Feigl (1961). "Why Ordinary Language Needs Reforming". *The Journal of Philosophy 58*(18), 488–498.

Mulhall, Stephen (1999). *Stanley Cavell: Philosophy's Recounting of the Ordinary*. Oxford: Oxford University Press.

Murphy, Taylor S. (2015). "Experimental Philosophy: 1935-1965". In T. Lombrozo, J. Knobe, and S. Nichols (Eds.), *Oxford Studies in Experimental Philosophy*, Volume 1, pp. 325–368. Oxford: Oxford University Press.

Myers, James (2009). "Syntactic Judgment Experiments". *Language and Linguistics Compass 3*(1), 406–423.

Naess, Arne (1938). *"Truth" As Conceived by Those Who Are Not Professional Philosophers*. Oslo: Kommisjon Hos Jacob Dybwad.

Nichols, Shaun, N. Ángel Pinillos, and R. Mallon (2016). "Ambiguous Reference". *Mind 125*(497), 145–175.

Plunkett, David and T. Sundell (2013). "Disagreement and the Semantics of Normative and Evaluative Terms". *Philosophers' Imprint 13*(23), 1–37.

Plunkett, David and T. Sundell (2014). "Antipositivist Arguments from Legal Thought and Talk". In G. Hubb and D. Lind (Eds.), *Pragmatism, Law, and Language*, pp. 56–75. Routledge.

Potter, Stephen (2005). *Lifemanship: Some Notes on Lifemanship with a Summary of Recent Research in Gamesmanship*. London: Moyer Bell and its subsidiaries.

Ryle, Gilbert (1949). *The Concept of Mind* (2009 ed.). New York, NY: Routledge.

Ryle, Gilbert (1953). "Ordinary Language". *The Philosophical Review 62*(2), 167–186.

Sandis, Constantine (2010). "The Experimental Turn and Ordinary Language". *Essays in Philosophy 11*(2), 181–196.

Schlenker, Philippe (2003). "Indexicality, Logophoricity, and Plural Pronouns". In J. Lecarme (Ed.), *Research in Afroasiatic Grammar II*, Selected Papers from the Fifth Conference on Afroasiatic Languages, pp. 409–428. Amsterdam: Benjamins.

Searle, John (1969). *Speech Acts*. Cambridge: Cambridge University Press.

Sprouse, Jon (2013). "Acceptability Judgments". *Oxford Bibliographies* (http://www.oxfordbibliographies.com/view/document/obo-9780199772810/obo-9780199772810-0097.xml).

Sprouse, Jon, C. T. Schütze, and D. Almeida (2013). "A Comparison of Informal and Formal Acceptability Judgments Using a Random Sample from *Linguistic Inquiry* 2001–2010". *Lingua 134*, 219–248.

Sundell, Timothy (2010). "Disagreements about Taste". *Philosophical Studies 155*(2), 267–288.

Tennessen, Herman (1962). "Permissible and Impermissible Locutions". In *Logic and Language*, Volume 5 of *Synthese Library*, pp. 220–233. Dordrecht: Springer.

Vendler, Zeno (1967). "Linguistics and the *A Priori*". In *Linguistics in Philosophy*, pp. 1–32. Ithaca, NY: Cornell University Press.