

Linguistic Corpora and Ordinary Language: On the Dispute between Ryle and Austin about the Use of ‘Voluntary’, ‘Involuntary’, ‘Voluntarily’, and ‘Involuntarily’¹

Michael Zahorec, Robert Bishop, Nat Hansen, John Schwenkler, and Justin Sytsma

Abstract

The fact that Gilbert Ryle and J.L. Austin seem to disagree about the ordinary use of words such as ‘voluntary’, ‘involuntary’, ‘voluntarily’, and ‘involuntarily’ has been taken to cast doubt on the methods of ordinary language philosophy. As Benson Mates puts the worry, ‘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). In this chapter, we evaluate Mates’s criticism alongside Ryle’s and Austin’s specific claims about the ordinary use of these words, assessing these claims against actual examples of ordinary use drawn from the British National Corpus (BNC). Our evaluation consists in applying a combination of methods: first aggregating judgments about a large set of samples drawn from the corpus, and then using a clustering algorithm to uncover connections between different types of use. In applying these methods, we show where and to what extent Ryle’s and Austin’s accounts of the use of the target terms are accurate as well as where they miss important aspects of ordinary use, and we demonstrate the usefulness of this new combination of methods. At the heart of our approach is a commitment to the idea that systematically looking at actual uses of expressions is an essential component of any approach to ordinary language philosophy.

(1) Introduction

In the middle of the 20th century, some philosophers argued that reflection on the language of non-philosophers could be a way of making progress on perennial philosophical problems. This approach, which came to be called ‘ordinary language philosophy’, pointed to supposed differences between the way philosophers used expressions that play central roles in philosophical arguments (expressions such as ‘looks’ in arguments for sense data, or ‘knows’ in arguments for skepticism) and the way those expressions are used outside philosophy. Failing to pay sufficient attention to ordinary use of expressions got philosophers into trouble, the ordinary language philosophers argued, because philosophers ended up using those expressions in ways that were so different from their ordinary uses that they in effect were either giving these expressions entirely new meanings (cf. Waismann 1997, Fischer 2019) or using them without

¹ Forthcoming in *Experimental Philosophy of Language: Perspectives, Methods and Prospects*, D. Bordonaba (Ed.), Springer. We want to thank David Bordonaba, Eugen Fischer, and Kevin Reuter for their helpful comments. Nat Hansen gratefully acknowledges support from the Alexander von Humboldt Foundation, John Schwenkler from the Alexander von Humboldt Foundation and the Notre Dame Institute for Advanced Study.

any real meaning at all (cf. Baz 2017), and so in any case weren't really talking about *looking* or *knowing* (or whatever topic they took themselves to be investigating) in any recognizable sense. Further, by distorting the meaning of expressions like 'looks' and 'knows', philosophers got themselves into unnecessary and unfruitful philosophical entanglements.² The way out of those entanglements was to reflect on the ordinary uses of the relevant expressions.

During the 1950s and early 1960s, ordinary language philosophy was arguably 'the most influential school of philosophy in Britain' (Russell 1953, p. 303), but it also attracted intense criticism from several different directions. The best-known objection to ordinary language philosophy holds that while philosophers may be departing from ordinary *use* when they employ expressions like 'looks' and 'know' in philosophical arguments, they aren't thereby distorting the *meaning* of those expressions (Grice 1961, Stroud 1984). A lesser known, but even more fundamental, objection criticizes what should be the core strength of ordinary language philosophy: its claims about the ordinary use of terms and expressions. If ordinary language philosophers fail to correctly describe the way that expressions are ordinarily used, then they have stumbled at the very first hurdle in their criticism of the way philosophers use language. This is what Mates (1958) argued in an early debate with Cavell (1958). Mates held that there is reason to doubt the accuracy of ordinary language philosophers' claims because Gilbert Ryle and J. L. Austin, two of the most influential and astute ordinary language philosophers, disagree in some of their claims about ordinary use. The conclusion that Mates draws from Austin and Ryle's disagreement is posed as a (rhetorical) question: '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?' (1958, p. 165).

The purported disagreement that Mates finds between Ryle and Austin concerns the expressions 'voluntary', 'involuntary', 'voluntarily', and 'involuntarily'. Mates observes that, in *The Concept of Mind*, Ryle makes the following claims:

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. (Ryle 1949/2009, p. 56)

Ryle goes on to say, in accordance with the general strategy pursued by ordinary language philosophers, that philosophers apply these terms in 'quite another way' from the ordinary use, namely as applying to not only actions that ought not to be done, but also to 'meritorious actions' (ibid.). It is this, he argues, that leads philosophers into nonsense and confusion.

² For a contemporary version of this argument, see Fischer, Engelhardt, Horvath, & Ohtani (2021).

Mates then points out that Ryle's remarks seem inconsistent with something Austin says about 'voluntarily' and 'involuntarily' in the context of a discussion of expressions we use to make excuses in ordinary language:

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)

Whatever one's views about the morality of joining the army, it seems clear that making a gift, hiccoughing or making a small gesture are not usually seen as 'actions which ought not to be done'. It therefore looks like Austin has given an example that shows that Ryle's characterization of ordinary use is incorrect. And this is the way the apparent conflict between Austin and Ryle has been characterized in the literature on ordinary language philosophy in the 60 years since Mates made his argument: Austin's example reveals Ryle's mistake (e.g., Cavell 1958, p. 174; Hacker 1996, p. 235; Hanfling 2000, p. 56; Norris 2017, p. 30).

We want to linger over this apparent disagreement. We say 'apparent' because, as any acute observer of ordinary language should notice, Austin's examples don't actually contradict what Ryle says in the passage that Mates quotes. This is because Ryle makes a claim only about the adjectives 'voluntary' and 'involuntary', while Austin's claim is about the adverbs 'voluntarily' and 'involuntarily'. However, not only does all of the existing commentary that addresses Mates's objection take it for granted (without argument) that Austin's and Ryle's claims are in conflict,³ but all existing commentary that we are aware of also fails to notice that, just a few pages later in *The Concept of Mind*, Ryle himself describes the use of 'voluntarily' in a way that is similar to Austin—right down to the example of volunteer soldiers:

Very often we oppose things done voluntarily to things suffered under compulsion. Some soldiers are volunteers, others are conscripts; some yachtsmen go out to sea voluntarily, others are carried out to sea by the wind and tide. (Ryle 1949/2009, p. 60)

Following this remark Ryle makes several further claims about different possible uses that 'voluntarily' and 'involuntarily' have in ordinary language, which we will discuss in detail below. It seems plausible, then, that Mates was simply wrong that there was disagreement between Ryle and Austin about the ordinary use of these expressions.⁴

Nevertheless, while the way that Mates frames the disagreement between Ryle and Austin is sloppy, it is possible to reconstruct a modified version of his challenge. In the passage quoted above, Ryle says that 'very often we oppose things done voluntarily to things suffered under compulsion', while Austin's position is that 'voluntarily' and 'involuntarily' can only be used to

³ One exception is Hansen (2017), who does note that Ryle and Austin's claims are not strictly speaking in conflict with each other.

⁴ For further discussion of this point, see Schwenkler (forthcoming).

modify normal verbs when the action named by the verb is done ‘in some special way or circumstances’:

The natural economy of language dictates that for the standard case covered by any normal verb,—not, perhaps, a verb of omen such as ‘murder,’ but a verb like ‘eat’ or ‘kick’ or ‘croquet’—no modifying expression is required or even permissible. Only if we do the action named in some special way or circumstances, different from those in which such an act is naturally done (and of course both the normal and the abnormal differ according to what verb in particular is in question) is a modifying expression called for, or even in order... It is bedtime, I am alone, I yawn: but I do not yawn involuntarily (or voluntarily!)... To yawn in any such peculiar way is just not to just yawn. (Austin 1957, p. 16)

At the root of Mates’s challenge is the following question: *How could we tell* whether Ryle or Austin (or both, or neither) is correct in their characterizations of ordinary language? In this chapter we will demonstrate the value to ordinary language philosophy of looking at a sample of actual language drawn from a *linguistic corpus*—a searchable body of text, purpose-built for answering linguistic questions (Bluhm 2016, p. 91).⁵

The results of our analysis pull in several different directions. First, our results cast doubt on several of the specific claims that Ryle and Austin make about the ordinary use of our target terms. This reveals the pitfalls of doing ordinary language philosophy without a systematic survey of the varieties of actual use. Second, our results nevertheless vindicate, in unexpected ways, some more general aspects of Ryle’s and Austin’s claims about these expressions, including Austin’s insistence that the pairs ‘voluntary’ and ‘involuntary’, and ‘voluntarily’ and ‘involuntarily’, do not function as simple opposites, and Ryle’s claim these words are used to mark a number of different non-overlapping conceptual distinctions. Third, we find that some characteristically philosophical uses of these expressions are closest in their use to one narrow type of ordinary use, namely one that categorizes bodily movements in physiological terms. This last finding anticipates a claim made by Anscombe (1963, p. 12) and raises the possibility of a novel response to Ryle and Austin’s claim that philosophical uses of these expressions depart from their ordinary use, namely that the philosophical use is actually continuous with *one* type of ordinary use. However, as we will discuss in Sect. 5, the philosophical significance of the similarity between philosophical and physiological uses of these expressions remains an open question.

⁵ Experimental philosophers have increasingly been calling on tools from corpus linguistics in recent years, treating corpora as a further source of evidence that can help with testing philosophical hypotheses about language. See Liao and Hansen (2022), Hansen et al. (2021), Sytsma et al. (2019), Caton (2020), and Ulatowski et al. (2020) for recent discussions and examples.

(2) Ryle on ‘Voluntary’ and ‘Involuntary’

Our first study began by considering the above-quoted passage from *The Concept of Mind*, focusing on the claims highlighted here in italics:

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*. (Ryle 1949, p. 56; emphasis added)

Ryle makes three claims in this passage about how ‘voluntary’ and ‘involuntary’ are ordinarily used: namely that (‘in their most ordinary employment’ and ‘with a few minor elasticities’) these words are used to (1) *describe actions* that (2) *ought not to be done*, (3) in cases where *the action seems to have been the fault of the agent*. To aid in a qualitative assessment of the accuracy of these claims, we used the British National Corpus (BNC) to generate a sample of 100 uses of each of the terms ‘voluntary’ and ‘involuntary’, together with their surrounding context.⁶ These two lists of 100 ‘key words in context’ (KWICs) comprised our sample corpora for this study, as used them to evaluate Ryle’s three claims by posing the following three questions of each KWIC:

- (Q1) Is the term of interest used to describe an action or some actions?
- (Q2) Does the speaker say, suggest, or assume that the agent(s) (either potentially or actually) ought not to have performed the action(s)?
- (Q3) Does the speaker say, suggest, or assume that the agent(s) are (either potentially or actually) at fault for something?

⁶ The BNC is a 100-million-word sample of British English from the late 20th century. 90% of the corpus is drawn from written materials, including newspapers, fiction and non-fiction books, and letters, while the remaining 10% comes from spoken language, including transcripts of government meetings, trials, and radio and television shows. In the BNC, there are 3849 uses of ‘voluntary’ and 359 uses of ‘involuntary’. To obtain our sample, we downloaded the entire BNC and used a straightforward text-sorting algorithm to isolate all the uses of each term of interest, along with the context in which each use appeared. With the help of a random number generator, we selected 100 entries from each of these lists. Subsequent examination revealed that two items out of the 200 were duplicates (both concerned uses of ‘involuntary’). These were removed from the analysis. The remaining entries, together with those from the study reported in the next section, are given in full in the supplemental materials. These are numbered by order in our original samples, with the items for ‘voluntary’ running from 1-100 and the items for ‘involuntary’ from 101-200. Shortened versions are shown in Figures 3-4 below. Examples of usage taken from the British National Corpus (BNC) were obtained under the terms of the BNC End User License. Copyright in the individual texts cited resides with the original IPR holders. For information and licensing conditions relating to the BNC, please see the web site at <http://www.natcorp.ox.ac.uk/>




Figure 1: Histograms of responses for each of the three questions in our first study (Q1-Q3), broken down by term (‘voluntary’ top, ‘involuntary’ bottom). Total responses are 500 for ‘voluntary’ (five coders for each of the 100 entries) and 490 for ‘involuntary’ (five coders for each of the 98 entries remaining after removing duplicates).

This study followed a simple procedure. Using Qualtrics, we presented each KWIC individually to each of the authors of this paper. Each author then answered each of Q1, Q2, and Q3 on a six-point scale with the response options ‘Not Applicable,’ ‘No,’ ‘Probably Not,’ ‘Not sure,’ ‘Probably,’ and ‘Yes,’ in reference to the passage in question. All five of the authors completed the survey before the responses were analyzed. Interrater reliability was fair to moderate across the questions and terms, suggesting that collectively we had some difficulty interpreting Ryle’s and applying them to real world examples—a point that we return to below.⁷

⁷ Interrater reliability was measured in two ways. First, coding ‘NA’ responses as 0, ‘No’ responses as 1, and so on, the results were treated as interval on a 0-6 scale. Intraclass correlation coefficients with 95% confidence intervals were then calculated for each term and for each question, treating both the items and raters as random effects. Results for ‘voluntary’ were: (Q1) 0.67 [0.59, 0.74], (Q2) 0.53 [0.44, 0.62], (Q3) 0.53 [0.44, 0.62]. Results for ‘involuntary’ were: (Q1) 0.44 [0.34, 0.54], (Q2) 0.52 [0.43, 0.61], (Q3) 0.54 [0.45, 0.63]. Second, we treated responses as categorical, combined negative responses (‘No’ or ‘Probably Not’) and positive responses (‘Yes’ or ‘Probably’). We then calculated Fleiss’s Kappa for each term and for each question. Results for ‘voluntary’ were: (Q1) 0.55, (Q2) 0.52, (Q3) 0.52. Results for ‘involuntary’ were: (Q1) 0.18, (Q2) 0.22, (Q3) 0.21.

Histograms of responses for each question are shown in Figure 1, and full results are given in the supplemental materials along with the full text of each KWIC we evaluated, indexed by target expression and position in the sample.

The first things to note from our results are the high number of items that generated negative responses (either ‘No’ or ‘Probably Not’) to Q1, especially for uses of ‘voluntary’, as well as the correspondingly high number of ‘NA’ responses to Q2 and Q3—both questions that presume we are dealing with an action. This suggests that, contrary to Ryle’s initial assumption, these terms are fairly often used to describe events that are not actions. Indeed, we gave predominantly negative responses (i.e., plurality ‘No’ or ‘Probably not’ with less than half ‘Probably’ or ‘Yes’) to Q1 on 53 of the 100 items for ‘voluntary’ and 23 of the 98 items for ‘involuntary’. While it is unclear just what frequency of use might be countenanced as a ‘few minor elasticities’, each proportion is greater than and significantly different from a conservative noise threshold of 15%.⁸ Looking at the individual items that generated these responses, a few groups stand out. One large group comprises items where ‘voluntary’ is used to say that certain ‘bodies’, ‘sectors’, ‘foundations’, or ‘organizations’ are neither for-profit nor part of the government, such as in the following:

[57] The committee includes representatives of local authorities, health authorities, Government departments and the *voluntary* and private sector.

Likewise, an example of a use of ‘involuntary’ that doesn’t apply to an action is the following, in which it is used to name that part of an animal’s nervous system that controls the muscles and glands of its internal organs:

[165] Her work showed that the cells migrate to many different sites in the embryo developing into the skeletal elements of the head, all the pigment cells in the body, most of the nerves of the *involuntary* nervous systems, sensory nerves, and a variety of glands.

Given that Ryle’s concern is only with the use of words like ‘voluntary’ and ‘involuntary’ in characterizing actions, we chose to exclude the items with predominantly negative responses to Q1 from the subsequent analysis.

However, even after restricting the sample to uses of ‘voluntary’ and ‘involuntary’ that were judged to describe actions, Ryle’s other two specific claims about the ordinary use of ‘voluntary’ were still not supported by our findings. In fact, a majority of the authors answered ‘No’ or ‘Probably Not’ to *both* Q2 and Q3 for *all but one* of the remaining 47 items.⁹ The one exception, for which the authors’ ratings were ambivalent, is still not a clear example of the use of ‘voluntary’ that Ryle seems to have in mind. That example concerns ‘voluntary liquidation’, a

⁸ ‘voluntary’: $\chi^2=110.3$, $p<.001$, 95% CI [0.43, 0.63]; ‘involuntary’: $\chi^2=4.87$, $p=.027$, 95% CI [0.16, 0.33].

⁹ Needless to say, this proportion is above and significantly different from the 15% threshold: $\chi^2=246.7$, $p<.001$, 95% CI [0.87, 1.00].

process by which the assets of an insolvent entity (person or company) are sold off to pay creditors:

[20] For another hour the inquisition continued, almost, I felt, as though the judges were scraping for any dirt they could find. One member asked for details of how my father had gone into *voluntary* liquidation.

Strikingly, many of the examples of ‘voluntary’ in the sample are used in precisely the circumstances which Ryle suggests the term won’t be used—for the word often modifies actions which *ought* to be done and for which *praise* is called for, as in the following:

[67] When it emerges that she has been doing her *voluntary* work in York for just six years, her enormous commitment becomes clear. Pets as Therapy is, of course, nothing new. But Joan ventures where few other dog owners would dare to tread...

Other uses characterize actions that are evaluatively neutral, such as:

[88] Furthermore, the clinical outcome seemed also to be related to the length of prolonged *voluntary* anal contraction achieved by patients.

Overall, based on our sample, it appears that all of Ryle’s three claims about the ‘most ordinary employment’ of ‘voluntary’ are mistaken, at least on a reasonable interpretation of a ‘few minor elasticities’ and with regard to contemporary use.

Ryle’s claims about the ordinary use of ‘involuntary’ fared slightly better. Of the 75 uses of ‘involuntary’ that we regarded as modifying actions, a majority of the authors answered ‘No’ or ‘Probably Not’ to *both* Q2 and Q3 for 44 of them.¹⁰ Indeed, there were a total of just six items that received positive responses from each of the authors for both Q2 and Q3, and another five that received majority positive responses to both questions. Five of the six items with unanimously positive responses employed the phrase ‘involuntary manslaughter’ ([106, 117, 119, 125, 126]), while the other expressed a worry from ‘Right-to-Life groups’ concerning the slippery slope leading toward ‘active and involuntary euthanasia’ ([194]). Further, many of the other items that received a majority positive response to Q2 and Q3 also seem to fit Ryle’s claims. For instance, the following sentence is naturally read as an excuse or apology for encroaching on someone’s private space:

[111] My intrusion was *involuntary*, Dr Vaughan. I can’t have meant to come here if it’s private property.

¹⁰ This proportion is, once again, significantly greater than the 15% threshold employed above: $\chi^2=108.8$, $p<.001$, 95% CI [0.47, 0.70].

Similarly, [189] describes an ‘involuntary remark’ that was ‘swiftly regretted’, while [198] details a question that ‘was totally involuntary’, noting that the speaker ‘could have bitten her tongue out for asking it’.

Nevertheless, a clear majority of uses of ‘involuntary’ in our sample do not align with Ryle’s account. Indeed, as noted above, a majority of the items received a negative response from the majority of the authors for both Q2 and Q3, as in the following examples:

[120] All are striking, some are beautiful, others startling and a few may invoke an *involuntary* shudder ...

[135] Her thoughts turned to Geoffrey Howe, for so long her most faithful lieutenant, and her right leg made an *involuntary* kicking movement.

While each of these sentences might be read as implying that the agent did not explicitly want to do the thing described, none of them seem to require that the agent *ought* not have done this, nor that it is something for which they should be held *at fault*. Once again, Ryle’s claims about the ordinary use of ‘involuntary’ simply do not square with how the word is most often employed in our sample.

(3) Further Claims: Austin’s ‘Special Circumstances’ and Beyond the Received View of Ryle

As we discussed in Sect. 1, the received view about Ryle’s account of the ordinary use of ‘voluntary’ and ‘involuntary’ is that it is simply mistaken. The results reported in the previous section look like clear support for this. But, as we discussed, the standard criticism fails to notice that Ryle goes on to make several further claims about this family of terms that go beyond what is said in the better-known passage about ‘voluntary’ and ‘involuntary’ that was the focus of Sect. 2. And Austin, like Ryle, also makes claims about the use of ‘voluntarily’ and ‘involuntarily’ that stand in need of assessment. This is what we did in our second study, deriving a new set of five questions from Austin’s and Ryle’s further remarks and assessing them in reference to the sample KWICs generated in our first study for the terms ‘voluntary’ and ‘involuntary’, as well as two further corpora, compiled from the BNC in the same way as before, of 100 uses each of the terms ‘voluntarily’ and ‘involuntarily’.¹¹

¹¹ The BNC contains 474 uses of ‘voluntarily’ and 214 uses of ‘involuntarily’. Of the 100 examples of each that were selected, six were found to be duplicates (one for ‘voluntarily’ and five for ‘involuntarily’). These were removed from the analysis. The remaining entries are given in full in the supplemental materials, numbered by order in our original sample (‘voluntary’ items numbered from 201-300, ‘involuntary’ items from 301-400), and shortened versions are shown in Figures 5-6 below.

Our first new question was derived from Austin. Above we quoted his different diagnosis of where the philosophical use of the modifying adverbs ‘voluntarily’ and ‘involuntarily’ goes awry. For Austin, the use of such a modifying expression is ‘required or even permissible’ only if the action it describes is done ‘in some special way or circumstances, different from those in which such an act is naturally done’ (Austin 1957, p. 16). To assess this claim we used the following question:

- (Q4) In using the term of interest, does the speaker say, imply, or suggest that the agent(s) did the thing in question in some special way or circumstances, different from those in which such an act is naturally done?

The remaining four questions in our second study derive from Ryle. Above we cited a neglected passage in *The Concept of Mind* where Ryle seems to anticipate the possibility of using a word like ‘voluntarily’ to describe an act, like becoming a soldier or going out to sea, that is not obviously something that ought not to be done nor that is someone’s fault. Here is the wider context of that remark, with a few crucial passages highlighted in italics:

Very often we oppose things done voluntarily to things suffered under compulsion. Some soldiers are volunteers, others are conscripts; some yachtsmen go out to sea voluntarily, others are carried out to sea by the wind and tide ... What is involuntary, in this use, is not describable as an act. Being carried out to sea, or being called up, is something that happens to a person, not something which he does ... So sometimes the question ‘Voluntary or involuntary?’ means ‘Did the person do it, or was it done to him?’; sometimes it presupposes that he did it, but means ‘Did he do it with or without heeding what he was doing?’ or ‘Did he do it on purpose or inadvertently, mechanically, or instinctively, etc.?’ (Ryle 1949, p. 60; italics added)

The most striking feature of this passage is that Ryle seems to be identifying several *different*, and possibly non-overlapping, ways that our terms of interest can be used. We took this as a cue to ask the following further questions:

- (Q5) Does the speaker use the term of interest as a way of clarifying whether the agent(s) did the thing in question on purpose, in contrast with doing it merely by accident?
- (Q6) Does the speaker use the term of interest as a way of clarifying whether the agent(s) did the thing in question of their own accord, i.e., not under threat or compulsion (by another person or natural forces, for example)?

- (Q7) Does the speaker use the term of interest as a way of clarifying whether the agent(s) really did the thing in question at all, rather than its being something that merely happened to them?
- (Q8) Does the speaker use the term of interest as a way of clarifying whether the agent(s) did the thing in question heeding what they were doing, as opposed to doing it inadvertently, mechanically, or instinctively?

It is worth pointing out that Q5-Q8 all ask whether the speaker is clarifying matters in regard to some particular contrast (e.g., ‘on purpose’ versus ‘merely by accident’ in Q5). Accordingly, in contrast with Q2-Q4, positive responses to these questions do not differentiate between cases where the speaker is using the term to say what falls on the first side of the contrast (e.g., that the action was done on purpose) and cases where the speaker is using the term to say what falls on the second side of the contrast (e.g., that the action was done by accident).

Following the method of our first study, in our second study each of the authors answered these five questions for all the items in our samples of uses of ‘voluntary’ and ‘involuntary’, as well as for the sample uses of ‘voluntarily’ and ‘involuntarily’. As in our first study, each KWIC was presented individually via Qualtrics and the questions were answered using the same six-point scale. Interrater reliability was generally just fair to moderate, and often poor for Q4 and Q7, again indicating the difficulty of interpreting these criteria and applying them to actual examples.¹² Histograms of responses for each question are shown in Figure 2, and breakdowns of responses for each item are given in the supplemental materials.

As with Q2 and Q3 from the previous study, our new questions all assume that we’re dealing with actions. It therefore isn’t surprising that we again found a high percentage of ‘NA’ responses for ‘voluntary’ (46.0%), and to a lesser extent ‘involuntary’ (9.3%), across our responses to these new questions. These responses were not evenly distributed, however. Across the two terms, 69 items (out of 198) had more than 30% ‘NA’ responses, accounting for 88.6% of the ‘NA’ responses overall. And, as expected, these were predominantly the items we classified as non-actions in Study 1 (88.4%). By contrast, ‘NA’ responses were rare for ‘voluntarily’ (0.4%) and ‘involuntarily’ (0.5%), suggesting that these terms almost always modify actions. (Given that these are adverbs, this is not surprising.) For the subsequent analysis,

¹² Interrater reliability was measured in the same two ways for our first study. ‘voluntary’: (Q4) ICC=0.45 [0.36, 0.55], Kappa=0.16; (Q5) ICC=0.53 [0.44, 0.62], Kappa=0.35; (Q6) ICC=0.68 [0.60, 0.75], Kappa=0.33; (Q7) ICC=0.54 [0.45, 0.63], Kappa=0.32; (Q8) ICC=0.63 [0.54, 0.71], Kappa=0.36. ‘involuntary’: (Q4) ICC=0.35 [0.26, 0.46], Kappa=0.12; (Q5) ICC=0.44 [0.35, 0.54], Kappa=0.22; (Q6) ICC=0.60 [0.51, 0.68], Kappa=0.50; (Q7) ICC=0.27 [0.18, 0.37], Kappa=0.11; (Q8) ICC=0.69 [0.61, 0.76], Kappa=0.47. ‘voluntarily’: (Q4) ICC=0.22 [0.13, 0.31], Kappa=0.08; (Q5) ICC=0.43 [0.34, 0.53], Kappa=0.28; (Q6) ICC=0.47 [0.38, 0.57], Kappa=0.23; (Q7) ICC=0.16 [0.08, 0.26], Kappa=0.06; (Q8) ICC=0.45 [0.36, 0.55], Kappa=0.25. ‘involuntarily’: (Q4) ICC=0.44 [0.34, 0.54], Kappa=0.25; (Q5) ICC=0.26 [0.17, 0.37], Kappa=0.05; (Q6) ICC=0.70 [0.63, 0.77], Kappa=0.58; (Q7) ICC=0.18 [0.10, 0.28], Kappa=0.09; (Q8) ICC=0.74 [0.67, 0.80], Kappa=0.50.

we restricted our samples to items that were classified as actions according to Q1 in our first study and then received less than one-third ‘NA’ responses to Q4-8.¹³

Our first question in the present study is Q4, which concerns Austin’s claim about ‘voluntarily’ and ‘involuntarily’ being used only to describe things that are done in ‘some special way or circumstances’. Only 20.2% of the uses of ‘voluntarily’ and 46.3% of the uses of ‘involuntarily’ in our sample were judged by a majority to satisfy Austin’s claim. Thus, the majority of the items for each term were not clearly used in the way that Austin claims is needed for the ‘modifying expression [to be] called for, or even in order’.¹⁴ Here are two examples that did seem to match Austin’s description, each receiving unanimous affirmative answers to Q4:

[281] And the only examination that I ever *voluntarily* and with malice aforethought ... failed ... was the Bank examination

[321] Once more he raised his arm *involuntarily*, as if in greeting.

Arguably, these two sentences are examples in which the terms of interest are used in reference to acts that were done in some special way or circumstances—namely, failing an exam on purpose rather than, as is typical, doing so despite attempting to pass; or raising one’s arm, but not in the ‘usual’ way of waving hello.

Alongside these Austin-friendly examples, however, many more of the occurrences in our sample used our terms of interest to describe actions that did not seem to be done in a special way or circumstances, given the kinds of actions that they are. (As Austin writes, ‘both the normal and the abnormal differ according to what verb in particular is in question’ (Austin 1957, p. 16).) Consider *shuddering*, *shivering*, or *trembling*, for example:

[343] Mait shuddered *involuntarily*.

[350] She shivered *involuntarily*, a reaction prompted by something other than cold.

[398] Quite *involuntarily* Isabel began to tremble.

We will consider examples like these in more detail below in discussing Q8, but the thing to emphasize for present purposes is that what ‘involuntarily’ seems to be indicating in these examples is only that the act in question was a mere bodily response (that is, that it was done ‘inadvertently, mechanically, or instinctively’, as Ryle puts it). And, of course, this is exactly what is *typically* the case for acts like shuddering, shivering, and trembling. These therefore look like

¹³ This left us with 40 items for ‘voluntary’ and 74 items for ‘involuntary’. Other than the duplicate items that were removed as described in Footnote 6, no further items were removed for ‘voluntarily’ or ‘involuntarily’, leaving us with the original set of 99 unique items for the former and 95 for the latter.

¹⁴ Both proportions are, of course, above and significantly different from the 15% threshold used previously: $\chi^2=321.0$, $p<.001$, 95% CI [0.70, 0.87]; $\chi^2=108.5$, $p<.001$, 95% CI [0.43, 0.64].

counterexamples to Austin’s claim that ‘involuntarily’ is only used to describe actions done in ‘some special way or circumstances, different from those in which *such* an act is normally done’.

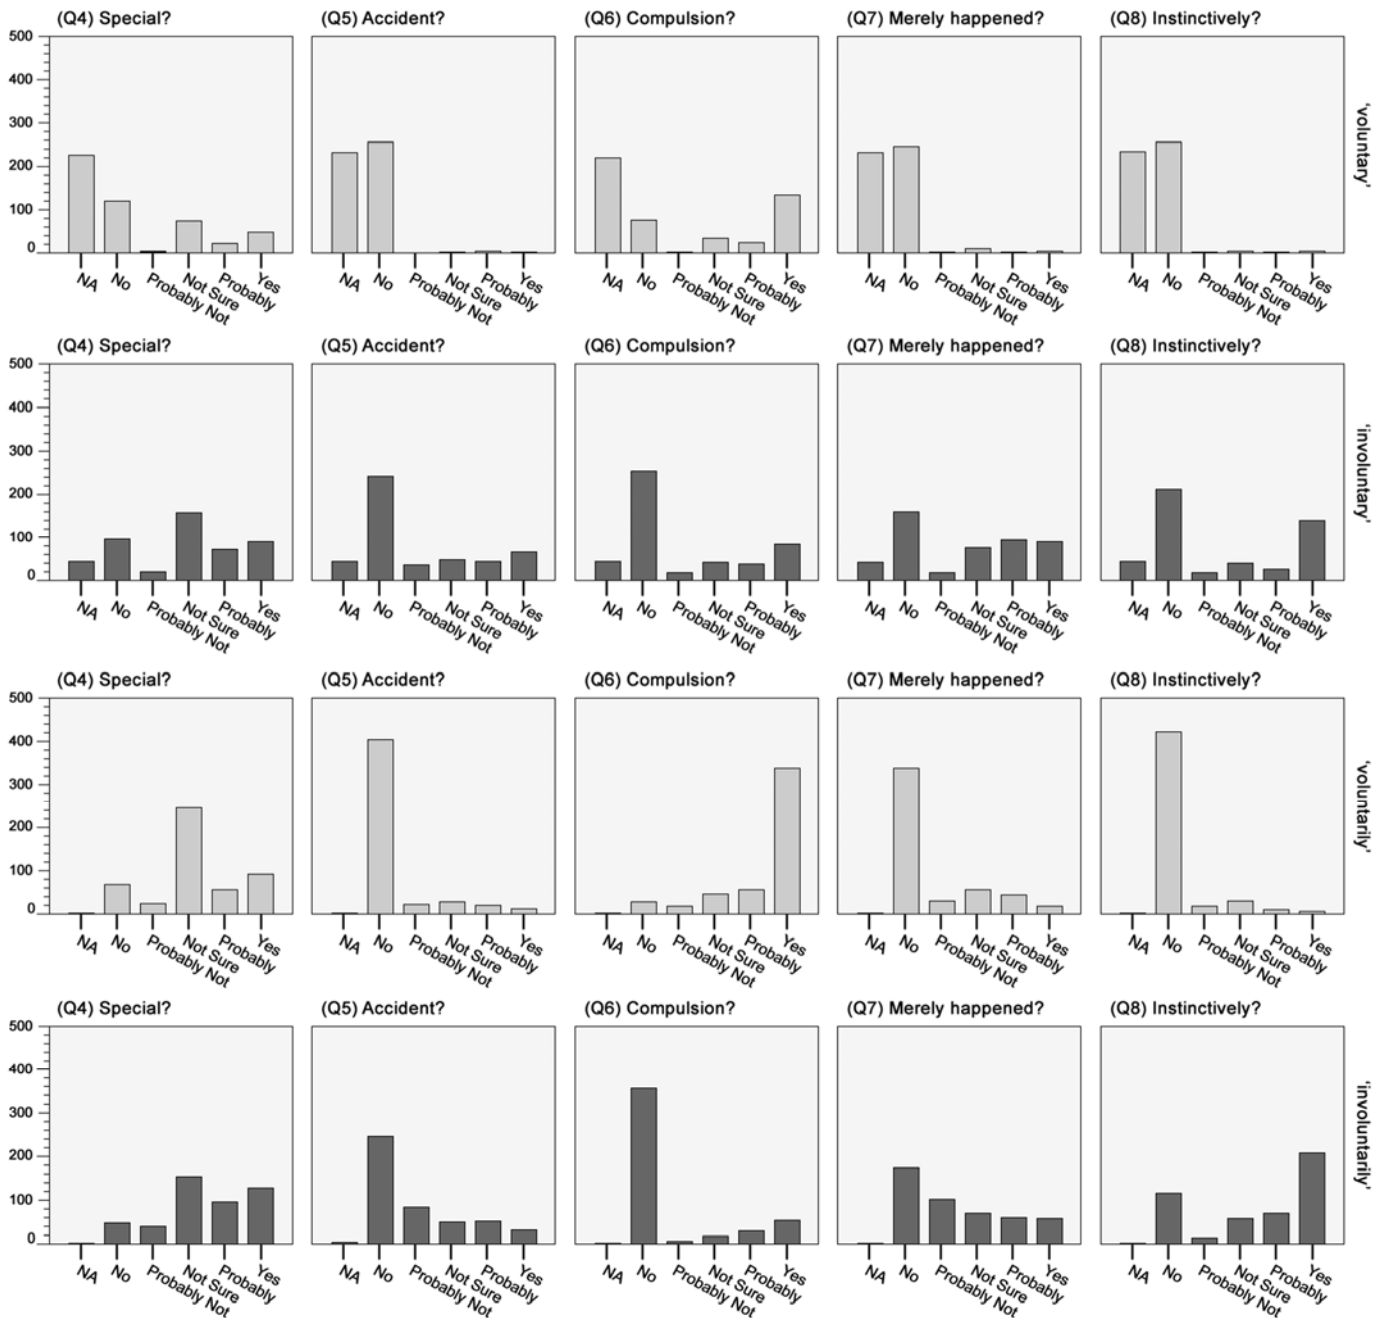


Figure 2: Histograms of responses for each of the five questions in our second study (Q4-Q8), broken down by term (from top to bottom: ‘voluntary’, ‘involuntary’, ‘voluntarily’, ‘involuntarily’). Total responses are 500 for ‘voluntary’ (five coders for each of the 100 entries), 490 for ‘involuntary’ (five coders for each of the 98 entries remaining after removing duplicates), 495 for ‘voluntarily’ (five coders for each of the 99 non-duplicate entries), and 475 for ‘involuntarily’ (five coders for each of the 95 non-duplicate entries).

Still, charity suggests thinking further about these cases, as there is an argument to be made that these uses of ‘involuntarily’ are not simply redundant or inelegant, but instead are being used to convey the idea that the bodily response in question stems from some unconscious fear or desire.¹⁵ If this is correct, then perhaps these aren’t counterexamples to Austin’s claim after all: for example, the ‘normal’ way to tremble is in response to something fearful, rather than to underlying romantic desire. Yet this only leads to a more serious problem for Austin’s claim, namely that often it is totally unclear what counts as a ‘special way or circumstance’, given the difficulty of saying what counts as the *normal* way of, or circumstances for, doing the kind of thing in question. Consider, for example, the following:

[213] Maj. Gen. Rodrigo Sanchez Casillas became the new Army Chief of Staff and Brig.-Gen. Garin Aguirre became the new Army Inspector General both replacing those who had ‘*voluntarily* resigned’ over the La Cutufa affair.

What are the normal conditions for an officer to resign? One might expect that sometimes officers resign under pressure or out of necessity, and on other occasions they do so freely and without undue pressure, perhaps to take a private sector job or make a political statement. But since it’s not obvious which of these possibilities should be counted as ‘special’, it is hard to say how Q4 should be answered in relation to it.

The difficulty of applying Austin’s description to examples of ‘voluntarily’ and ‘involuntarily’ was a common one. Indeed, nearly half of our responses to Q4 for ‘voluntarily’ (49.9%) and a third of our responses for ‘involuntarily’ (32.8%) were ‘Not Sure’. Not surprisingly, across the two terms only a minority of items (32.8%) secured majority agreement to Q4, and even fewer found us agreeing unanimously (8.2%). And although Austin only mentions this principle with regard to the adverbs ‘voluntarily’ and ‘involuntarily’, it is worth noting that we encountered the same difficulty with ‘voluntary’ and ‘involuntary’ as well: 31.8% of the judgments for these items were ‘Not Sure’, and just 28.1% of items secured majority agreement (2.6% unanimous).

The rest of the questions that we posed in this second study are drawn from Ryle’s distinctions between different ways that our terms of interest can be used. In contrast with the claim that we drew from Austin, Ryle’s claims generated distinctions that were more readily applied to the sample. Thus, while we were very often uncertain in our responses to Q4, this happened far less frequently for the remaining four questions, as we chose the response ‘Not sure’ just 8.9% of the time overall. Not only were we generally able to apply Ryle’s distinctions, but each seemed to clearly apply to some of the uses in our sample. Nonetheless, as we will see, these distinctions don’t apply equally clearly to each use, nor do they collectively exhaust the variety of ways the terms of interest are used—indeed, far from it.

¹⁵ Specifically, in [343] the shudder is a response to a remark about being close to open water, in [350] the shiver is a response to the unexpected silence of an empty castle, and in [398] the trembling is a response to the ‘low laugh’ of a seducer.

As is clear from Figure 2, the questions about compulsion (Q6) and whether something was done inadvertently, mechanically, or instinctively (Q8) produced the most positive responses. A majority of us responded ‘Yes’ or ‘Probably’ to Q6 for 47.7% of items and to Q8 for 31.8% of items, but Q5 (on purpose vs. by accident) and Q7 (something really ‘done’ at all?) received majority positive responses for only 8.1% and 14.9% of the items, respectively. Interestingly, however, our positive responses to Q6 and Q8 were associated with different sets of terms, as Q6 tended to be answered positively in connection with the use of ‘voluntary’ and ‘voluntarily’, while for Q8 the positive responses were primarily for uses of ‘involuntary’ and ‘involuntarily’ instead. Specifically, we found that while a majority of us responded positively to Q6 for 67.5% of the uses of ‘voluntary’ and 85.9% of the uses of ‘voluntarily’, a majority of us gave positive responses for only 24.0% of uses of ‘involuntary’ and 17.9% of uses of ‘involuntarily’. The inverse was found for Q8, as a majority of us responded positively for just 2.5% of uses of ‘voluntary’ and 1.0% of uses of ‘voluntarily’, compared to 46.7% of uses of ‘involuntary’ and 64.2% of uses of ‘involuntarily’. As we discuss in Sect. 5, this finding anticipates a point that Austin makes about the use of the adverbs ‘voluntarily’ and ‘involuntarily’: that in their ordinary use these words ‘are not opposed in the obvious sort of way that they are made to be in philosophy or jurisprudence’, and so ‘in spite of their apparent connexion, are fish from very different kettles’ (1957, p. 17).

With regard to the question of compulsion, Ryle draws the contrast, as framed in Q6, between things that are *done* and things that are *suffered* under compulsion, suggesting that we sometimes use ‘voluntarily’ (and ‘voluntary’) to mark the first category, and ‘involuntarily’ (and ‘involuntary’) to mark the second. And our study did turn up a number of examples that were judged to fit this characterization. For instance, consider the following uses of ‘voluntarily’:

[252] ... the Cabinet approved a plan which called on all groups and individuals *voluntarily* to hand over small-calibre weapons by mid-March.

[260] ... members dissatisfied with their union can *voluntarily* resign without the threat of losing their job.

In each of these passages, the use of the target term seems to describe an agent as having done something of their own accord in a context where we might otherwise assume that this was done only due to coercion or the threat of retaliation. Interestingly, for many of the items where ‘voluntary’ and ‘voluntarily’ were used in this way, the contrast with coercion seems to carry with it a suggestion of pressure, sometimes even serving as a direct threat—i.e., that if the agent does not do the thing in question ‘voluntarily’ then the desired outcome will be brought about in some other way. This is clearly the case for [252] above, which goes on to say that ‘the government threatened that after the deadline it would launch a campaign forcibly to collect the weapons, and warned that the “severest of penalties” would be inflicted on those violating the order’.

On the other side of the distinction, we also find many examples of ‘involuntarily’ and ‘involuntary’ being used to mark things that a person suffered, often through having something done to them by force. For instance, [129] describes the Politburo’s ouster of Khrushchev as ‘the first and only *involuntary* departure of an established Kremlin leader’, while [315] describes the commitment of the English poet Christopher Smart to St. Luke’s Hospital for Lunatics in terms of his ‘entering the asylum *involuntarily*’, and [345] concerns a member of a merchant vessel crew who was pressed into service, so that he ‘found himself *involuntarily* a member of the Royal Navy’. In each case, the suggestion is not simply that some pressure or coercion was brought to bear on a person in a way that influenced their choice of action; rather, ‘involuntarily’ is used to indicate that something was forcibly done to them: Khrushchev was removed, Smart imprisoned, Lyell impressed.

In this context, it is worth revisiting Q7, which was intended to track Ryle’s appeal to the contrast between ‘something that happens to a person, not something which he does’. While our initial expectation was that answers to this question would largely align with answers to Q6, in fact this is not what we found: a majority of us answered both Q6 and Q7 positively for only 11 total items, or 3.6% of our sample. Instead, responses to Q7 were much more aligned with those to Q8, which concerns the distinction between doing something ‘while heeding what [one was] doing, as opposed to doing it inadvertently, mechanically, or instinctively’. In particular, many of the items that generated positive responses to Q7 concerned bodily acts like shivering, shuddering, and trembling, while others concerned things that (in an ordinary sense of this phrase) were clearly done *to* people, often by a government agency or bureaucracy and described in a way that minimized responsibility. Examples include discussions of involuntary repatriation [148, 175] and involuntary reception into care [133, 153]. Indeed, only one item seemed to squarely fit with Ryle’s example of the yachtsmen carried out to sea by the wind and tide:

[319] The pilot stated that as the aircraft rose above the treeline, at about 150 feet above the ground it *involuntarily* banked to the right, and despite maintaining the climb speed he could not prevent the roll to the right which continued past ninety degrees of bank.

This suggests that our target terms are not generally used to specifically mark whether something *merely* happened to a person, but that they are much more frequently employed with regard to either whether something was done to a person or whether a person did something but without conscious intent.

Turning finally to Q8—which contrasts voluntary action with things that are done inadvertently, mechanically, or instinctively—our results are essentially the reverse of what we saw for Q6: while the terms ‘involuntary’ and ‘involuntarily’ are very often used to draw the distinction that is framed in Q8, the terms ‘voluntary’ and ‘voluntarily’ are used in this way much less often. Indeed, only two items in our samples for ‘voluntary’ and ‘voluntarily’ received majority positive responses to Q8. The first is an item, noted earlier, which describes medical research on the use of biofeedback to treat incontinence:

[88] Furthermore, the clinical outcome seemed also to be related to the length of prolonged *voluntary* anal contraction achieved by patients.

The second describes work on hypnosis, emphasizing that ‘the hypnotised subject is not a will-less automaton’:

[264] A subject may take up the hypnotic suggestion that he is unable to bend his arm: He is actively, deliberately, *voluntarily* keeping his elbow stiff while simultaneously orchestrating for himself the illusion that he is really trying his best to bend it.

The target expressions in these examples are plausibly intended to emphasize that the agent was doing the thing in question deliberately or on purpose—contrary, as Austin might have added, to how it is typically done (as in the first item), or to how it would have appeared to a naive observer (as in the second).

In items [88] and [264], the point of using ‘voluntary’ and ‘voluntarily’ seems to be to say that the act in question was not merely a physiological response or incapacity, but rather something that fell within the agent’s control. Many of the uses of ‘involuntary’ that generated favorable responses to Q8 seemed to concern a similar distinction, such as describing shivering as ‘a form of involuntary muscular action [that] raises the metabolic rate and elevates body temperature’ in [101] or speaking of ‘uncontrollable involuntary bodily movements’ in [155]. Further, numerous uses of ‘involuntarily’ also seemed to evoke this distinction—among them, the shivering described in [350] that was unrelated to the cold, the description in [312] of a man’s ‘hands clenching involuntarily’ after a shock, the discussion in [394] of the range of stress symptoms resulting when ‘muscles are involuntarily clenched’, and the following first-person narrative:

[351] Coming down the slippery track, I stumble. *Involuntarily* I reach out my arm.

This last item describes a bodily movement where the contrasts drawn in Q6 and Q8 are both salient: the movement in question is performed ‘involuntarily’ insofar as it is inadvertent or mechanical, rather than deliberate or purposive. However, for many of the bodily movements at issue, this contrast is not nearly so salient. For instance, as noted above acts like shivering, shuddering, and trembling are not generally done on purpose: indeed, there is a sense in which ‘shivering on purpose’, for example, would not be to actually shiver, but to fake it.

In closing our discussion of the results for our second set of questions, we want to note two further reasons for thinking that Ryle’s characterization of the different ways of using our terms of interest is less than fully adequate. First, we found that 10.4% of our items failed to elicit majority favorable responses to *any* of the four questions based on his remarks, suggesting that they involve uses of our terms of interest that Ryle’s characterizations fail to capture. While this included items for each of the four terms, the largest subset was for ‘voluntary’, with 30% of the actional uses

failing to receive a majority favorable response to any of our questions. These items included multiple instances of the phrases ‘voluntary workers’ ([46, 47]) and ‘voluntary work’ ([67]), as well as ‘voluntary contributions’ ([10]), ‘voluntary dog walkers’ ([30]), ‘voluntary helpers’ ([82]), and ‘voluntary assistance’ ([87]), among related uses. In these items, the point of describing the acts in question as ‘voluntary’ seems to be to say that they are things that the agents in question *volunteered* to do, as opposed to doing this for pay or as part of an official capacity. Ryle’s description of the many different ways that ‘voluntary’ can be used leaves out this possibility.

Second, within those items that did elicit majority positive responses to any given question, there was often notable heterogeneity in what the terms of interest seemed to be used to say. Consider for example the following four sentences, each of which received mainly positive responses to Q5:

[104] Tics are involuntary movements. Like Martial and Rabelais, Mozart’s lavatory humour in his letters, poems, and canons (for example, *Leck mich im Arsch*: lick my arse, K231) was not *involuntary* but intentional.

[374] Beetles had fed on the pollen of cycads and they were among the first to transfer their attentions to the early flowers like those of magnolias and waterlilies. As they moved from one to another, they collected meals of pollen and paid for them by becoming covered in excess pollen which they *involuntarily* delivered to the next flower they visited.

[88] Furthermore, the clinical outcome seemed also to be related to the length of prolonged *voluntary* anal contraction achieved by patients.

[253] Taking these several points in combination we come to the particular significance of the shedding of blood in the ritual of circumcision. The belief was that to ritually and voluntarily—and I stress the word *voluntarily*—shed one’s own blood was to recommend oneself to and establish a link with the Creator of the Universe, and this is precisely what happened with circumcision.

There are important differences between these examples. In [104] and [88] a contrast is drawn between an involuntary movement and an action that involves the agent’s conscious control, but in [374], the involuntary action (pollination) is something that is done inadvertently, as a by-product of the activity of feeding on pollen, while in [253] the contrast is between freely choosing to do something and being coerced into doing it by some outside power (choosing to shed one’s own blood freely versus ‘suffering’ the procedure or being coerced into it). As such, it seems that the distinction drawn in Q5 was insufficient on its own to distinguish between these different uses, as it cross-cut some other important differences in how our terms of interest are used. The penultimate section of our chapter presents a further analysis that we conducted in order to see whether further differences like these could be brought out by looking at our responses taken in

aggregate, in order thereby to provide a more comprehensive and systematic overview of the ways that our terms of interest are ordinarily used.

(4) Categorizing Uses

In looking at actual examples of how people use the terms ‘voluntary’, ‘involuntary’, ‘voluntarily’, and ‘involuntarily’ in the past two sections, we’ve seen that while Ryle’s and Austin’s remarks pick up on some facets of the ordinary use of these terms, they are neither wholly accurate nor do they tell the full story. In many instances we found uses that were plausibly related to distinctions they drew, but nevertheless did not fit them squarely. In others we found uses that weren’t clearly related to these distinctions at all. Finally, in yet other cases we found ourselves unsure of how to respond to various items, or in disagreement amongst ourselves in how we did respond. In order to help reveal some deeper order in these response patterns, our final study used the technique of cluster analysis to group the items in our samples together based on our responses in our first two studies.

Cluster analysis includes a broad range of statistical procedures that aim to group items together in a way that minimizes the differences between items in a group (or ‘cluster’) and maximizes the differences between groups, all with regard to some relevant set of measurements.¹⁶ For our purposes, the items we wanted to cluster are our sample KWICs for ‘voluntary’, ‘involuntary’, ‘voluntarily’, and ‘involuntarily’, and the measurements are the mean responses to the questions that were presented in our first two studies. To minimize ‘NA’ responses, we used the restricted set of items discussed at the start of Sect. 3, with any remaining ‘NA’ responses excluded from the calculation of the means. Given the scope of the restricted set of items, we excluded Q1 from the analysis of ‘voluntary’ and ‘involuntary’, and given the close correlation between Q2 and Q3, these questions were combined into a single dimension. Further, given the difficulties with applying Q4 that we described earlier, this question was excluded for all terms. For the analysis we employed agglomerative hierarchical clustering using Ward’s method with Euclidean distance.¹⁷ To aid the interpretation process, cluster dendrograms were generated separately for each of our four terms of interest. The dendrograms that were generated by this process, overlaid with the conceptual distinctions that we identified, are shown in Figures 3-6.

¹⁶ For examples of the use of cluster analysis in experimental philosophy, as well as further discussion of the method, see Levine et al. (2021), Fischer and Sytsma (2021), Fischer, Engelhardt, and Herbelot (forthcoming), Sytsma and Snater (forthcoming), and Woike et al. (2020). See Reuter et al. (2022) for another example of the use of cluster analysis in looking at linguistic corpora.

¹⁷ See Sytsma and Snater (forthcoming) for a more detailed description of the procedure. As described there, if the analysis is being used for hypothesis testing, it is good practice to compare multiple clustering methods to test robustness. Given that our goal here was instead exploratory, aiming merely to help organize the items in a way that would be fruitful, we instead took a ‘proof is in the pudding’ approach, and as we’ll see our default setting produced intelligible clusterings of KWICs.

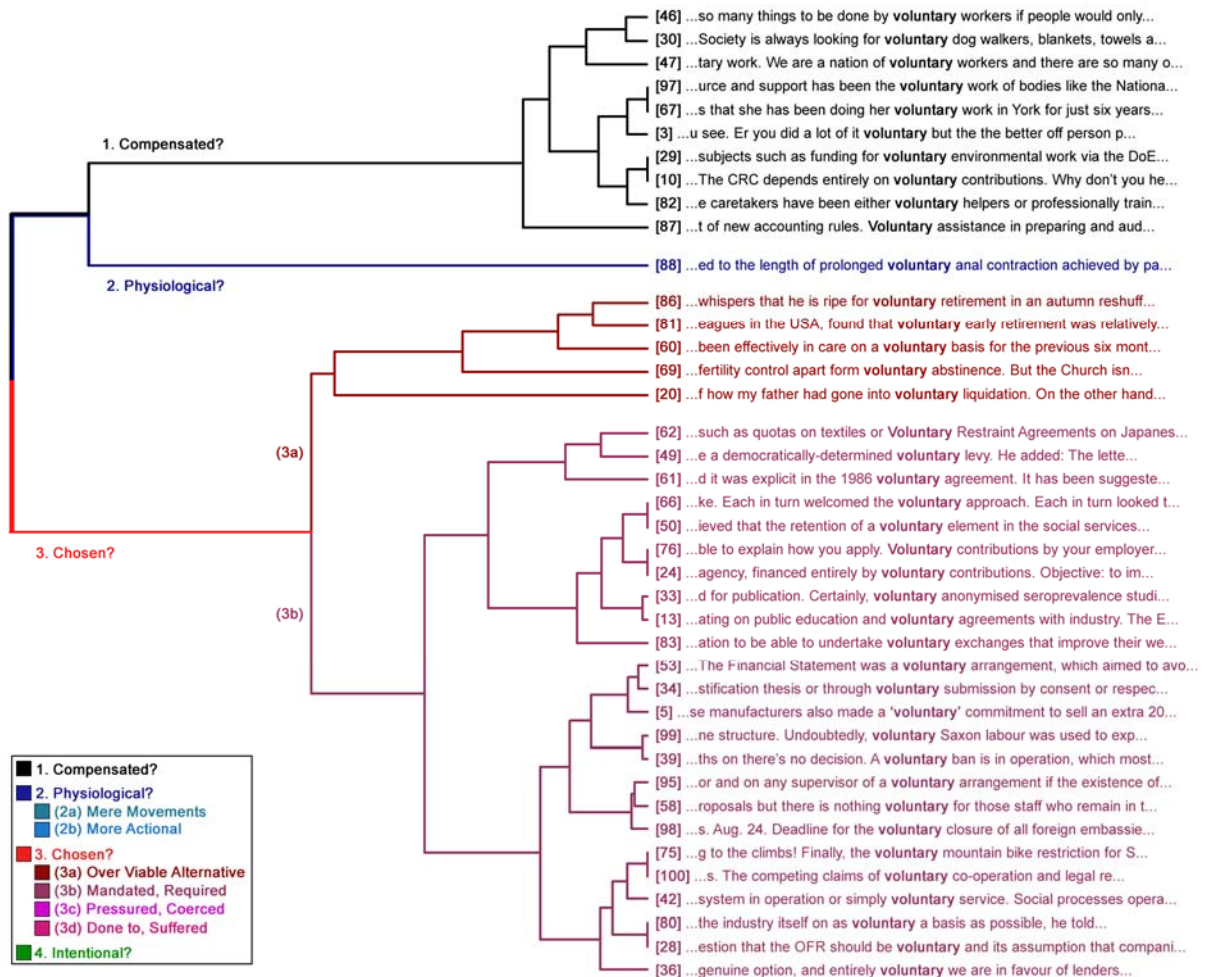


Figure 3: Annotated dendrogram for agglomerative hierarchical clustering using Ward's method with Euclidean distance for the restricted set of 40 items for 'voluntary' (items classified as actions in our first study and receiving less than one-third 'NA' responses in our second study). Items are labeled with their number from the original sample and abbreviated KWIC. Cluster coloring and labels corresponds with our interpretation of the conceptual distinctions captured by the cluster analysis for each of our four target terms, showing sub-distinctions where they could be discerned.




Figure 4: Annotated dendrogram for agglomerative hierarchical clustering using Ward's method with Euclidean distance for the restricted set of 74 items for 'involuntary' (items classified as actions in our first study and receiving less than one-third 'NA' responses in our second study). Items are labeled with their number from the original sample and abbreviated KWIC. Cluster coloring and labels corresponds with our interpretation of the conceptual distinctions captured by the cluster analysis for each of our four target terms, showing sub-distinctions where they could be discerned.

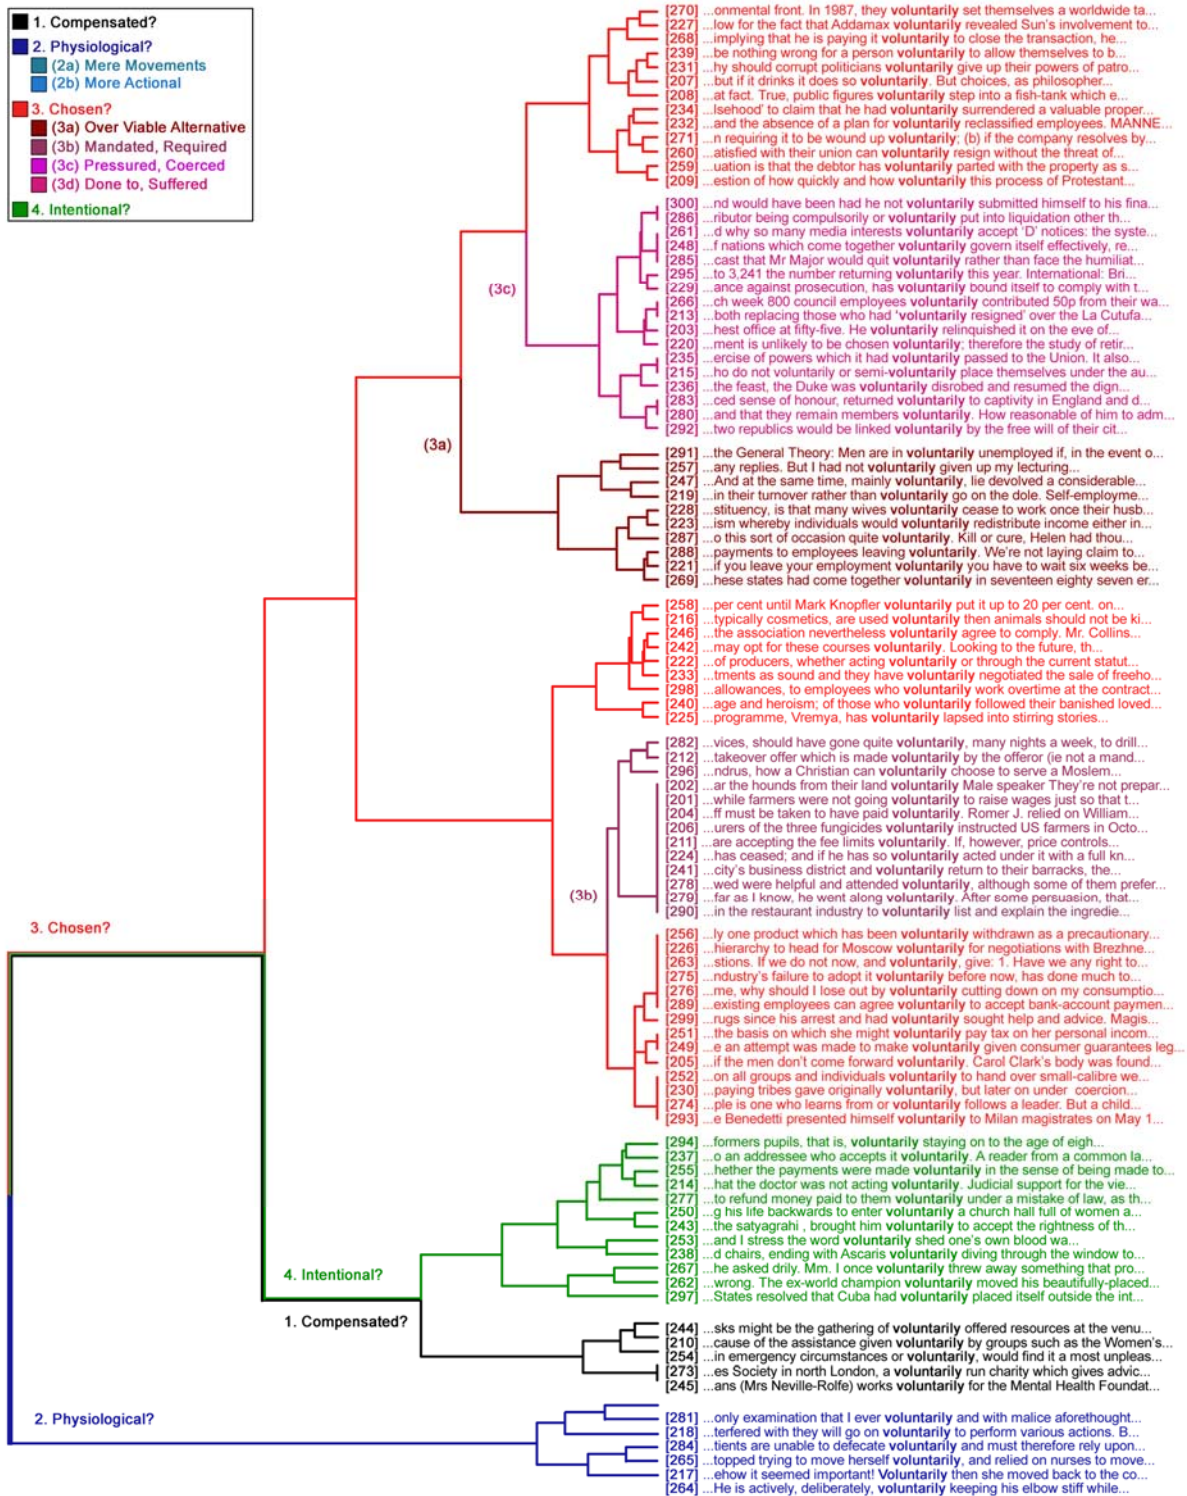


Figure 5: Annotated dendrogram for agglomerative hierarchical clustering using Ward's method with Euclidean distance for the 99 non-duplicate items for 'voluntarily'. Items are labeled with their number from the original sample and abbreviated KWIC. Cluster coloring and labels corresponds with our interpretation of the conceptual distinctions captured by the cluster analysis for each of our four target terms, showing sub-distinctions where they could be discerned.

- 1. Compensated?
- 2. Physiological?
 - (2a) Mere Movements
 - (2b) More Actional
- 3. Chosen?
 - (3a) Over Viable Alternative
 - (3b) Mandated, Required
 - (3c) Pressured, Coerced
 - (3d) Done to, Suffered
- 4. Intentional?




Figure 6: Annotated dendrogram for agglomerative hierarchical clustering using Ward’s method with Euclidean distance for the 95 non-duplicate items for ‘involuntarily’. Items are labeled with their number from the original sample and abbreviated KWIC. Cluster coloring and labels corresponds with our interpretation of the conceptual distinctions captured by the cluster analysis for each of our four target terms, showing sub-distinctions where they could be discerned.

Our goal in performing the cluster analysis was to uncover some structure from the complicated variation in our judgments about the target expressions. To this end, after generating the dendrograms we went on to explore them qualitatively, looking for semantic unity within the different uses that were clustered together and offering our own interpretation of the conceptual distinctions that these clusters suggest. The analysis performed admirably in this capacity, highlighting both some broad differences in the use of our target terms as well as some more subtle distinctions. Consider, as a first example, the clusters labeled (1) in the dendrograms, which tend to concern examples of things that were done on a *volunteer* basis rather than for the sake of compensation or reward, including volunteer labor and charitable donations. This is the use that we identified at the end of Sect. 3 as being overlooked in Ryle's analysis, and a significant number of uses of the modifiers 'voluntary' and 'voluntarily', as seen in Figures 3 and 5, were found to employ it. Notably, however, we did not find any examples at all of this use in connection with 'involuntary' and 'involuntarily'—another point in favor of Austin's advice not to treat these pairs as simple opposites. (For example, the opposite of '*voluntary* assistance' [87] is not 'involuntary assistance', but *paid* assistance.)

The second family of uses we identified—those labeled (2) in the dendrograms—has some presence among each of the four terms we investigated, but appears far more commonly in connection with 'involuntary' and 'involuntarily' than 'voluntary' and 'voluntarily'. These uses have to do with what we have called a *physiological* notion of voluntariness, centering on whether an action or movement is in some way automatic rather than under conscious control. Within the clusters that exemplified this use we found a further distinction, not always sharp, between how 'actional' the behaviors are. Thus, in one sub-category of uses, identified in our clusters as (2a), we find that the term of interest modifies mere bodily movements, as in 'involuntary spasms' [190] or 'involuntary muscular action' [101], as well as cases of shivering, shuddering, or trembling as discussed above. By contrast, in the second sub-category, identified as (2b), the examples tend to involve what are more naturally described as actions, such as in an 'involuntary yell of alarm' [173] or 'involuntary sobs' [163]. As we discuss in Sect. 5, this is the use of our terms of interest that arguably corresponds most closely to the philosophical notion of voluntariness as the basis of debates about free will. We therefore found it notable that it appears in our sample in connection with a relatively small range of descriptions, and—as observed above—hardly at all in the use of 'voluntary' and 'voluntarily'.

The third family of uses we identified corresponds roughly to Ryle's distinction between 'things done voluntarily' and 'things suffered under compulsion', though we found this specific distinction to be more of an endpoint along a spectrum. Perhaps the furthest from Ryle's description are a set of cases in (3a) that include phrases like 'voluntarily unemployed' [291], 'involuntary childlessness' [195], and 'voluntary abstinence' [69], where there does not seem to be any implication that the actions in question might have been compelled. Rather, by our lights the purpose of our terms of interest in these cases is to specify whether the agent chose to act in the way in question despite having been able to do something else, where it might have been expected that they would choose this alternative instead. The second set of choices we

distinguished, in (3b), involves a reference to a codified prescriptive norm such as a rule or regulation. That is to say, these items typically concern whether the action in question was done because it was *mandated or required*—for example, describing a certain program as a ‘democratically-determined voluntary levy’ [49] or discussing whether the restaurant industry would ‘voluntarily list and explain’ the ingredients in their meals [290]. Often the items that fell under this subclass involve actions that we’d usually consider good, and they are most often described using ‘voluntary’ or ‘voluntarily’—in some cases as a way of highlighting that credit is deserved, in others as a way of noting that something *wouldn’t* be done unless further pressure was brought to bear (e.g., ‘farmers were not going voluntarily to raise wages’ [201]). Third, in (3c) we found items that seemed to make reference to some form of explicit pressure or coercion. Relatively few examples of this use were explicitly identified in our clusters, but we can think of this use as contrasting with the one just discussed. Indeed, we find an example of just such an explicit contrast, with a hypothetical payment being described as ‘voluntary if no improper pressure was brought to bear, and involuntary if it was’ [149]. Finally, we have a larger set of uses classified under (3d) that correspond with what Ryle calls the difference between *suffering* something, or having it *done to* one, rather than *doing* it in a strict sense: for example, the point of saying that a duke ‘was voluntarily disrobed’ [236] is to describe this as his own act rather than an act that he suffered, while ‘involuntary reception into care’ [153] means being *put* into care rather than going there on one’s own.

The last use of our terms of interest that we identified in our clusters—those labeled under (4) in our dendrograms—seem to center around whether the thing in question was done intentionally or not. The uses of ‘involuntary manslaughter’ tended to fall under this category, as was the description in [111] of an accidental intrusion. This use is made quite explicit in a legal discussion of whether certain payments ‘were made voluntarily in the sense of being made to close the transaction’ [255], and was also displayed in the recollection of having ‘voluntarily [thrown] away something that promised to be special, and very, very wonderful’ [267]. There are, of course, many points at which some of these uses shade off into other ones, and the items that we have highlighted in our figures do not all correspond exactly to a given use. Still, we are struck by the extent to which these distinctions in use are borne out in the results of our cluster analyses, which seem to us therefore to speak in favor of the utility of this method as a tool for the exploration of ordinary meaning

To help validate the insights we drew from the cluster analyses, we employed a further technique from corpus linguistics, looking at where our terms of interest were located in a semantic space built from another common, general-purpose corpora (Corpus of Contemporary American English). We used the best performing distributional semantic model from Sytsma et al. (2019) to check the nearest neighbors of our target terms in the semantic space—the terms that the model says are closest in meaning. The results were striking, with the most synonymous terms generated suggesting the dominant categories we arrived at. For instance, the nearest neighbor for ‘voluntary’ was ‘mandatory’, suggesting (3b) from our classification, which was the

largest identified cluster. Similarly, the nearest neighbor for ‘involuntarily’ was ‘reflexively’, suggesting (2) from our classification, which again was the largest identified cluster.¹⁸

(5) Conclusion

Our close look at the sample of ordinary use drawn from the BNC has taken us far beyond Mates’s diagnosis of the supposed disagreement between Ryle and Austin about the use of ‘voluntary’, ‘involuntary’, ‘voluntarily’, and ‘involuntarily’. But we have vindicated Mates’s general worry that ordinary language philosophy should not be practiced without a systematic survey of the way language is ordinarily used, as we have seen how both Ryle and Austin overlook some types of ordinary uses of these expressions and how that leads to incomplete general observations about the features of ordinary language. For example, Austin warns against treating ‘voluntary’ and ‘involuntary’ and ‘voluntarily’ and ‘involuntarily’ as simple opposites. While we have identified one specific instance in which that observation holds true—the opposite of ‘voluntary assistance’ [87] is not ‘involuntary assistance’, but assistance that is *paid*—it clearly does not apply to all examples of the ordinary use of these expressions: the opposite of ‘voluntary anal contraction’ [88] is ‘involuntary anal contraction’.

This *physiological* use of ‘voluntary’ and ‘involuntary’ is something both Ryle and Austin seem to fail to notice, or at least seem not to appreciate the significance of. Ryle and Austin’s neglect of this use was noted by G. E. M. Anscombe in *Intention*, where she argues that we should ‘reject a fashionable view of the terms “voluntary” and “involuntary” [according to which] they are appropriately used only when a person has done something untoward’, suggesting that anyone ‘tempted by this view ... should consider that physiologists are interested in voluntary action, and they are not giving a special technical sense to that word’ (Anscombe 1963, p. 12). As we have discussed, many of the sentences falling under category (2) in our dendrograms supply clear illustrations of the type of use that Anscombe highlights, revealing that this is a point at which Ryle and Austin miss the significance of one aspect of the ordinary use of these expressions.

¹⁸ See Sytsma et al. (2019) for a further explanation of distributional semantic models and the semantic space used. The twenty nearest neighbors for each term of interest in terms of cosine (given in parentheses) are as follows. ‘voluntary’: mandatory (0.73), mandated (0.67), mandate (0.64), voluntarily (0.63), compulsory (0.63), compliance (0.63), comply (0.60), require (0.60), discriminatory (0.59), prohibit (0.59), implement (0.59), licensure (0.59), exempt (0.58), restrictive (0.58), participation (0.58), encourage (0.58), participate (0.58), requirement (0.57), stringent (0.57), cessation (0.57); ‘involuntary’: involuntarily (0.69), convulsive (0.61), spasmodic (0.60), manslaughter (0.59), forcible (0.57), induce (0.55), spasm (0.55), uncontrollable (0.54), agonized (0.54), convulsion (0.53), nausea (0.53), voluntary (0.53), bodily (0.52), breathlessness (0.52), gagging (0.52), tetany (0.52), immobility (0.51), asphyxia (0.51), peristaltic (0.51), forced (0.51); ‘voluntarily’: refuse (0.74), obligate (0.67), consent (0.65), permission (0.64), willingly (0.64), voluntary (0.63), legally (0.63), lawfully (0.61), allow (0.61), permit (0.61), request (0.61), unwilling (0.58), reluctant (0.57), reluctantly (0.56), coerce (0.56), subsequently (0.56), comply (0.56), illegally (0.56), notify (0.56), tacitly (0.55); ‘involuntarily’: involuntary (0.69), reflexively (0.63), volition (0.58), immobilize (0.57), shudder (0.57), weakly (0.58), paralyzed (0.56), flinch (0.56), numb (0.56), revolted (0.56), spasm (0.55), convulse (0.55), uncontrollably (0.55), contort (0.55), numbness (0.54), violently (0.54), convulsive (0.54), unconscious (0.54), unconsciousness (0.53), immediately (0.53).

The fact that both Austin and Ryle overlook the physiological use of these expressions raises a more challenging question about the methodology of ordinary language philosophy, a question on which the authors of this chapter are divided: can the use of these words that Anscombe highlights be drawn on in a defense of traditional philosophical discussions of voluntary action? Consider the following examples of clearly *philosophical* uses of our terms of interest, both of which appeared in the cluster of physiological uses:

[107] What, then, is moral luck? Nagel observes that it is intuitively plausible that people cannot be morally assessed for what is not their fault, or for what is due to factors beyond their control. Of course this makes sense in the case of insanity, automatism or *involuntary* movement but the range of factors over which one has no control is obviously wider than such clear instances of total lack of control.

[144] Notice also that expressive behaviour in this definition does not distinguish between voluntary and involuntary behaviour. If I jump with surprise when a dog suddenly barks at me, my behaviour is no less expressive than if I shout at it to shut up.

It is, we think, significant that both of these uses of ‘involuntary’ in abstract philosophical discussions are clustered alongside physiological uses like ‘involuntary discharge of faeces and urine’ [142] and ‘involuntary muscular action’ [101]. Since, as Anscombe says, the latter use of ‘involuntary’ does not involve giving the word any ‘special technical sense’, then perhaps the same is true of the philosophical uses that clustered with it. If that’s right, then our survey of ordinary use provides the resources for a novel rebuttal that the traditional philosopher can give to the ordinary language philosopher: philosophical use is an extension of one particular sub-type of ordinary use, and it is meaningful to the extent that that sub-type of ordinary use is meaningful.

The authors of this chapter are divided over how successful they think this reply to the ordinary language philosopher is. Some of us find it significant that the operative way of using these words is infrequent (only one use in our entire sample for ‘voluntary’, none for ‘involuntarily’, and only a small group each for ‘involuntary’ and ‘voluntarily’), and also that it only appears in connection with a narrow range of descriptions (for discussion, see Schwenkler forthcoming). But this is not the place to settle that debate. The essential lesson we wish to draw is the methodological point that the only way to make progress on resolving this kind of debate is by following Wittgenstein’s (1953, §66) advice: if one wishes to see what is common and different among the ways that we use our words, then “don’t think, but look!”

References

- Anscombe, G. E. M. (1963). *Intention*. ([2d. ed.]). Cornell University Press.
- Austin, J. L. (1956). A Plea for Excuses: The Presidential Address. *Proceedings of the Aristotelian Society*, 57, 1–30.
- Baz, A. (2017). *The Crisis of Method in Analytic Philosophy*. Oxford University Press.
- BNC Consortium. (2007). *The British National Corpus, XML Edition*, Oxford Text Archive, <http://hdl.handle.net/20.500.12024/2554>.
- Bluhm, R. (2016). Corpus Analysis in Philosophy. In M. Hinton (Ed.), *Evidence, Experiment, and Argument in Linguistics and Philosophy of Language* (pp. 91–109). Peter Lang.
- Caton, J. (2020). Using Linguistic Corpora as a Philosophical Tool. *Metaphilosophy*, 51(1), 51–70.
- Cavell, S. (1958). Must we mean what we say? *Inquiry*, 1(1–4), 172–212.
- Davies, Mark. (2008-) *The Corpus of Contemporary American English (COCA)*. Available online at <https://www.english-corpora.org/coca/>
- Fischer, E. (2019). Linguistic legislation and psycholinguistic experiments: redeveloping Waismann’s approach. In D. Makovec & S. Shapiro (Eds.), *Friedrich Waismann: The Open Texture of Analytic Philosophy* (pp. 211–241). Palgrave Macmillan.
- Fischer, E., Engelhard, P., Horvath, J., & Ohtani, H. (2021). Experimental ordinary language philosophy: a cross-linguistic study of defeasible default inferences. *Synthese*, 198, 1029–1070.
- Fischer, E., & Sytsma, J. (2021). Zombie intuitions. *Cognition*, 215, 104807.
- Grice, H. P. (1961). The Causal Theory of Perception. *Proceedings of the Aristotelian Society, Supplementary Volumes*, 35, 121–152.
- Hacker, P. M. S. (1996). *Wittgenstein’s place in twentieth-century analytic philosophy*. Blackwell.
- Hanfling, O. (2003). *Philosophy and Ordinary Language: The Bent and Genius of our Tongue* (2nd ed.). Routledge.
- Hansen, N. (2017). Must We Measure What We Mean? *Inquiry* 60(8), 785–815.
- Hansen, N., Porter, J. D., & Francis, K. (2021). A Corpus Study of “Know”: On the Verification of Philosophers’ Frequency Claims about Language. *Episteme*, 18(2), 242–268.

- Levine, S., Rottman, J., Davis, T., O’Neil, E., Stich, S., & Edouard, M. (2021). Religious Affiliation and Conceptions of the Moral Domain. *Social Cognition*, 39(1), 139–165.
- Liao, S., & Hansen, N. (2022). “Extremely Racist” and “Incredibly Sexist”: An Empirical Response to the Charge of Conceptual Inflation. *Journal of the American Philosophical Association*, 1–23.
- Mates, B. (1958). On the verification of statements about ordinary language. *Inquiry*, 1(1–4), 161–171.
- Norris, A. (2017). *Becoming who we are: Politics and practical philosophy in the work of Stanley Cavell*. Oxford University Press.
- Reuter, K., Baumgartner, L., & Willemsen, P. (2022). Tracing Thick and Thin Concepts Through Corpora. <http://philsci-archive.pitt.edu/20584/>
- Russell, B. (1953). The Cult of “Common Usage.” *The British Journal for the Philosophy of Science*, 3(12), 303–307.
- Ryle, G. (2009). *The Concept of Mind: 60th Anniversary Edition*. Routledge.
- Schwenkler, J. (forthcoming). Knowledge of language as self-knowledge. *Inquiry*.
- Stroud, Barry. (1984). *The significance of philosophical scepticism*. Clarendon Press.
- Sytsma, J., Bluhm, R., Willemsen, P., & Reuter, K. (2019). Causal Attributions and Corpus Analysis. In E. Fischer & M. Curtis (Eds.), *Methodological Advances in Experimental Philosophy* (pp. 209–238). Bloomsbury.
- Sytsma, J., & Snater, M. (Forthcoming). Consciousness, Phenomenal Consciousness, and Free Will. In P. Henne & S. Murray (Eds.), *Advances in Experimental Philosophy of Action*. Bloomsbury.
- Ulatowski, J., Weijers, D., & Sytsma, J. (2020). Cognitive Science of Philosophy Symposium: Corpus Analysis. *The Brains Blog*. <https://philosophyofbrains.com/2020/12/15/cognitive-science-of-philosophy-symposium-corpus-analysis.aspx>
- Waismann, F. (1997). *The Principles of Linguistic Philosophy*. Springer.
- Wittgenstein, L. (2009). *Philosophical Investigations, 4th edition (trans. Hacker and Schulte)*. Wiley-Blackwell.
- Woike, J. K., Collard, P., & Hood, B. (2020). Putting your money where your self is: Connecting dimensions of closeness and theories of personal identity. *PLoS ONE*, 15(2).