

Ordinary Meaning and Consilience of Evidence¹

Justin Sytsma

Abstract: In this chapter I note two recent trends, one in experimental jurisprudence and one in experimental philosophy. First, some work in experimental jurisprudence has pushed for moving beyond textual sources, including the use of linguistic corpora, and toward questionnaire methods in analyzing ordinary meaning. Second, some work in experimental philosophy has urged that we should look to move beyond the use of questionnaire methods and toward the use of linguistic corpora in analyzing ordinary concepts. There is a methodological tension here that suggests further investigation. I do so by considering the legal hypothetical of a prohibition on vehicles in a park that has featured prominently in the back-and-forth over the use of corpus methods in legal interpretation. Taking a closer look at this example, including extending previous corpus analyses, I argue that corpus and questionnaire methods in fact paint a remarkably similar picture of the ordinary use of “vehicle.” I hold that this highlights how these methods can compliment each other, and conclude that when it comes to difficult empirical questions—such as those that arise in assessing ordinary meaning and ordinary concepts—we should aim to employ multiple sources of information to arrive at a consilience of evidence.

Both philosophy and legal interpretation are often concerned with assessing the ordinary meaning of important terms or phrases (or with the concepts beings expressed by those terms when used with their ordinary meaning).² And researchers in both areas have begun to urge that ordinary meaning is an empirical question. Interestingly, some recent work in these areas suggests seemingly opposed views about the use of one source of empirical evidence in investigating ordinary meaning—the tools of corpus linguistics. Some work in experimental jurisprudence has pushed for moving beyond textual sources, including the use of linguistic corpora, and toward survey methods in legal interpretation (Tobia 2020, Klapper et al. forthcoming). In fact, Klapper et al. seem to suggest that we should prefer survey methods *over*

¹ Forthcoming in *Advances in Experimental Philosophy of Law* edited by Stefan Magen and Karolina Prochownik. I want to thank Eugen Fischer, Kevin Tobia, and Pascale Willemsen for their extremely helpful suggestions.

² This is not to say that researchers in these areas are after exactly the same thing, and indeed legal interpretation typically focuses on the ordinary or public meaning of a term given the context of some legal text, while philosophers tend to focus on analyzing ordinary concepts more generally, even if this is sometimes restricted to certain contexts (such as that provided by a given thought experiment).

corpus methods, holding that “the best way to discern ordinary meaning is to simply ask ordinary people how they would read and apply a disputed term” and making the argument that “properly executed, survey methods offer marked theoretical advantages over current means of ascertaining ordinary meaning” (20). Such work leads two advocates of the use of corpus methods in legal interpretation to contend with “a range of scholars seeking to pump the brakes on or outright repudiate the utility of corpus tools in the law” (Lee and Mouritsen 2021: 278–9).

In contrast, questionnaire methods have been the dominant approach in experimental philosophy, although they are certainly not the only methods that have been employed (Sytsma and Livengood 2016). Recently, however, some philosophers have urged that corpus methods “should be added to the philosopher’s toolkit” (Caton 2020: 51) and that they can productively supplement the more common questionnaire methods employed by experimental philosophers.³ For instance, Bluhm (2016: 104) notes some potential issues with questionnaire methods and suggests that one way to alleviate them is to turn to corpus methods, “not to supplant, but to supplement.” In similar fashion, Sytsma et al. (2019: 233) conclude that corpus methods should be seen “as a fruitful addition to the methodological toolbox of experimental philosophy” and

³ And philosophers have increasingly been calling on corpus methods, ranging from targeted web searches, to the tools I’ll focus on here (frequency counts, collocates, concordance lines), to the creation of corpora and employment of sophisticated computational approaches. While the use of corpus methods in philosophy dates back to at least the 1970’s—e.g., Meunier et al. (1976), McKinnon (1970); thanks to Louis Chartrand for making me aware of this work—their use has expanded in recent years. A non-exhaustive list of recent, English-language work in philosophy employing or discussing corpus methods broadly construed includes: Ludlow (2005), Meunier et al. (2005), de Villiers et al. (2007), Knobe and Prinz (2008), Wright and Bartsch (2008), Reuter (2011), Sainte-Marie et al. (2011), Slingerland and Chudek (2011), Herbelot et al. (2012), Bluhm (2013, 2016), Nagel (2013, 2021), Overton (2013), Tallant (2013), Vetter (2014), Andow (2015a, 2015b), Fischer et al. (2015, forthcoming), Liao et al. (2016), Nichols et al. (2016), Wright et al. (2016), Pence (2016, forthcoming), Ramsey and Pence (2016), Allen et al. (2017), Fischer and Engelhardt (2017), Hahn et al. (2017), Murdock et al. (2017), Schwitzgebel and Dicey Jennings (2017), Sytsma and Reuter (2017), Alfano et al. (2018), Nichols and Pinillos (2018), van Wierst et al. (2018), Pence and Ramsey (2018), Alfano and Cheong (2019), Betti et al. (2019), Malaterre et al. (2019, 2020, forthcoming), Mejía-Ramos et al. (2019), Pease et al. (2019), Sytsma et al. (2019), Caton (2020), Hinton (2020), Mizrahi (2020a, 2020b, forthcoming-a, forthcoming-b), Weatherson (2020), Gastaldi (2021), Nichols (2021), Fischer and Sytsma (2021), Alfano (forthcoming), Allen and Murdock (forthcoming), Hansen et al. (forthcoming), Malaterre and Chartier (forthcoming), Nie (forthcoming), Sytsma (forthcoming-a), Tsugita et al. (forthcoming), Ginammi et al. (forthcoming), Bonino et al. (forthcoming), Transwell and Inglis (forthcoming), Lean et al. (forthcoming), Willemsen et al. (forthcoming), Sytsma and Snater (forthcoming), Mizrahi and Dickinson (2021).

Ulatowski et al. (2020) note that “one advantage of corpus methods for experimental philosophy is that they can offer a further way to test our hypotheses, one free of some encumbrances common in more standard experimental contexts, even as it inevitably introduces others.” The suggestion is that corpus methods can often complement more traditional experimental methods, with the hope of providing a consilience of evidence with regard to hypotheses about ordinary meaning and ordinary concepts.

At first glance, there would seem to be a tension in these attitudes about corpus methods. This seeming tension recommends taking a closer look, and doing so, I believe, helps to draw out several important methodological points concerning the use of corpus and questionnaire methods. Most importantly, I argue that there is seldom a silver bullet when it comes to answering complex empirical questions and that different types of methods, including corpus and questionnaire methods, can often be used together to paint a fuller picture of a phenomenon of interest. To draw this out I will focus on a famous legal hypothetical—the meaning of “vehicle” in ordinances like “no vehicles in the park” (Hart 1958)—that plays a prominent role in the recent methodological debates in legal interpretation.

Here is how I will proceed. In Section 1, I briefly discuss ordinary meaning and illustrate the corpus methods that I’ll focus on. In Section 2, I present and then expand upon Lee and Mouritsen’s corpus analysis of “vehicle.” In Section 3, I consider a specific concern raised by Tobia, the supposed nonappearance fallacy, and some of the questionnaire results he presents. In Section 4, I then consider the questionnaire results that Klapper et al. report.

1. Ordinary Meaning and Linguistic Context

Lee and Mouritsen (2018) advocate for the use of corpus methods in addressing empirical

questions about ordinary meaning in the law. One illustration they provide involves the iconic problem of the meaning of “vehicle” in the rule “no vehicles in the park” (Hart 1958). The same hypothetical is subsequently taken up by Tobia (2020) and Klapper et al. (forthcoming) who raise concerns about the use of the corpus methods employed by Lee and Mouritsen in assessing ordinary meaning. While these groups take opposing views with regard to how best to empirically assess the ordinary meaning of “vehicle,” employing either corpus methods or one of two types of questionnaire methods, I’ll argue that the results for these methods are in fact consistent and that together they provide a fuller picture of the ordinary meaning.

While Lee and Mouritsen, Tobia, and Klapper et al. all argue convincingly that ordinary meaning is of clear importance in legal interpretation, what is meant by “ordinary meaning” is not itself so clear. As Lee and Mouritsen (2018: 127) put it, “ironically, we have no ordinary meaning of ‘ordinary meaning.’” Given vagueness with regard to what we are after when we look to assess ordinary meaning, I suggest that a single type of study is unlikely to be satisfactory.

There are numerous issues to be disentangled here, including that different theories of legal interpretation will suggest different perspectives on “ordinary meaning.” One important issue concerns the role of linguistic context in shaping how a word or phrase is understood. It is clear that how a word or phrase is understood is often contingent on the context in which it is used. To give a simple illustration, asking someone to “set the table” suggests a quite different sense of “table” than if you were to ask them to “format the table.” And, in fact, these correspond with the first two senses of the noun “table” defined in a common online dictionary:

[1] A piece of furniture with a flat top and one or more legs, providing a level surface for eating, writing, or working at.

[2] A set of facts or figures systematically displayed, especially in columns. (Lexico.com, 26 April 2021)

With this in mind, one question we might ask with regard to ordinary meaning is how a word is most commonly used across contexts. And for such questions, the tools of corpus linguistics can be of great help. To illustrate, we can check a common, balanced corpus such as the *Corpus of Contemporary American English* (COCA) to get a sense of how often the noun “table” is used in something roughly like each of the dictionary senses just noted.⁴ For instance, continuing with our toy example we can look at the most common *collocates* of the noun “table”; that is, the terms that occur most frequently near this term in the corpus. The results for COCA suggest that the noun is most often used in the furniture sense, with the top four collocates each suggesting this sense (“sit,” “kitchen,” “coffee,” “chair”) and seven of the top 10 (“dinner,” “dining,” “round”). The remainder of the top 10 collocates, however, instead suggest the information display sense (“present,” “figure,” “content”).

Alternatively, we can use the *Key Word in Context* (KWIC) feature to provide a sample of the texts, or *concordance lines*, in which “table” is used as a noun in COCA. A survey of a random sample of 100 concordance lines suggests that a little more than half employ the furniture sense (or a use derived from it, such as phrases like “table manners” and “table cloth”), while roughly a third employ the information display sense (including references to database tables). Most of the remainder involve idiomatic or metaphorical expressions related to the furniture sense, such as “off the table” and “ran the table.” Together, these results suggest, not surprisingly, that “table” is most frequently used in the furniture sense, although the information display sense is also reasonably common.

⁴ <https://www.english-corpora.org/coca/>

We might also ask questions about the specifics of a given broad sense of a term. For example, we might ask just which things are typically counted as tables in the furniture sense and whether the dictionary definition given above adequately captures the extension of the term as ordinarily used. Again, the methods of corpus linguistics can help. For instance, looking through concordance lines for the noun “table” on COCA, we find uses such as “pool table” and “bedside table” that are not (typically or primarily) used for eating, writing, or working at. Noting such exceptions, we might ask about the relative frequency of talking about different types of tables (in the furniture sense). One reason to do so is that this will provide an initial indication of what people are most likely to assume is meant by a generic reference to a “table.”⁵ And here we might note, for instance, that the *frequency count* for occurrences of “dinner table” is roughly three times larger than for “pool table” and two and a half times larger than for “bedside table.”

With regard to legal interpretation, Lee and Mouritsen (2018: 795) note that “when we speak of *ordinary* meaning, we are asking an empirical question—about the sense of a word or phrase that is most likely implicated in a given linguistic context.” And they argue that the tools of corpus linguistics—focusing on the tools briefly introduced above (frequency counts, collocation data, and concordance lines)—can help answer such questions. With regard to our target example of “vehicle,” for instance, they argue that “corpus data can tell us the relative frequency of different senses of *vehicle*... in naturally occurring language” (829).

⁵ When encountering words, language users infer that their referents have the features stereotypically associated with the word (e.g., the “secretary” will be female), unless inferences are cancelled by the context (Levinson 2000). Use frequencies influence what inferences are made, at different levels. Stereotypical associations are built up not only through observation of co-occurrence frequencies in the physical environment, but also through exposure to co-occurrence patterns in linguistic usage (McRae and Jones 2013). Moreover, about 40% of English words have several distinct, but related senses (Byrd et al. 1987), and the inferences people draw from such polysemes are influenced by the relative use frequencies of their different senses. Where one sense is clearly dominant, comprehenders are liable to make inferences licensed only by that dominant sense also from subordinate uses (Fischer and Engelhardt 2017, 2020; Fischer et al. 2021).

The different senses of “vehicle” that they have in mind here, however, are *not* the broad senses that standard dictionaries aim to provide. Their concern is with ordinances such as “no vehicles in the park,” which provide a context that rather clearly excludes three of the four senses suggested by one online dictionary:

[1] A thing used for transporting people or goods, especially on land, such as a car, lorry, or cart.

[2] A thing used to express, embody, or fulfil something.

[3] A film, television programme, song, etc., that is intended to display the leading performer to the best advantage.

[4] A privately controlled company through which an individual or organization conducts a particular kind of business, especially investment. (Lexico.com, 26 April 2021)

The linguistic context of the ordinance makes clear that we’re concerned with something like the first sense. Correspondingly, Lee and Mouritsen’s concern in asking about the “relative frequency of different senses of *vehicle*” is with the range of things—the range of potential vehicles—that might reasonably be taken to fall under the broad “transporting” sense of the term and, relatedly, how well this extension is captured by standard dictionary definitions like [1]. Thus, Lee and Mouritsen (2018: 788) note that “when judges speak of ordinary meaning, they often seem to be speaking to a question of relative frequency,” where this is taken to fall along a continuum: “Whether we regard the ordinary meaning of a given word to be the *possible*, *common*, or the *most common* sense of that word in a given context, linguistic corpora allows us to determine empirically where a contested sense of a term falls” (831–832).

In other words, what seems to be at issue with regard to “sense” in Lee and Mouritsen’s discussion is just how broad the range of conveyances at issue is ordinarily taken to be in a given context. And while I believe that linguistic corpora can *help* with this task, I’m skeptical that corpus tools will provide a convincing answer on their own. One major reason for this is that

assessing the use of a word “in a given context” can be difficult with corpora, although the extent of the difficulty depends on just how finely grained we take the context to be. For instance, with regard to the ordinance “no vehicles in the park,” we might wonder whether the relevant context is discussions of parks, or prohibitions against vehicles, or prohibitions against vehicles in parks, or a specific prohibition against vehicles in a specific park. The more finely-grained the context, the less likely there are to be a suitable number of examples in a corpus to warrant drawing conclusions about typicality from an investigation of the corpus.

2. Corpus Analyses

Let’s turn to the specifics of Lee and Mouritsen’s illustrative corpus analysis for the contemporary usage of the term “vehicle.” As we’ve seen, they hold that “corpus data can tell us the relative frequency of different senses of *vehicle*” and they contend that “if the search for ordinary meaning entails analyzing the relative frequency of competing senses of a given term, then corpus linguistics seems the most promising tool” (829). To do this for contemporary usage, Lee and Mouritsen begin by presenting the fifty most common collocates of “vehicle” in the News on the Web corpus (NOW).⁶

Lee and Mouritsen find that the most common terms occurring in the vicinity of “vehicle” in the corpus include “electric,” “motor,” “cars,” “traffic,” and so on. They reasonably conclude that many of the collocates “strongly indicate *automobile* as a likely candidate for the most common use of the term.” They then go on to note that two types of conveyances that might plausibly be thought to fall under the category of “vehicle” are not found in the top collocates: “airplane does not appear, though two particular types of aircraft are attested in the

⁶ <https://www.english-corpora.org/now/>

collocates—unmanned aerial vehicles (drones) and spacecraft” and, similarly, that “*bicycle* does not appear among the collocates of *vehicle* in contemporary usage” (838). Lee and Mouritsen then suggest that this finding raises an important question: “if *vehicle* is never used to refer to *bicycle* or *airplane* in the corpus data, then we may end up with an even further extension of our frequency continuum from *possible but rare* to *possible but unattested*” (840).

Lee and Mouritsen note, however, that before concluding that the use of “vehicle” to refer to bicycles or airplanes is unattested, we should also evaluate the KWIC data for the term. To do this, they reviewed a random sample of 100 concordance lines for “vehicle” from the same corpus. What they found is that the vast majority of these were referring to automobiles (91%) and that “the NOW Corpus data included no *airplanes*, *bicycles*, *tricycles*, *skateboards*, *roller-skates*, *toy cars*, or any of what Hart and others have characterized as penumbral, disputed cases” (842). While Lee and Mouritsen do not draw the strong conclusion that current usage of “vehicle” does not include such penumbral cases, they do take the corpus results to be suggestive of that conclusion, writing that “some seemingly *possible* meanings are unattested and may not be current” (842).

This motivates Lee and Mouritsen to look further, searching for concordance lines including both “vehicle” and one or the other of two disputed terms—“bicycle” or “airplane.” They find a few texts in which bicycles are counted as vehicles, but report that they “were unable to find a single collocation or concordance line that reflected the use of *vehicle* to mean *airplane*” (844). And they take this to further raise the question of whether “*airplane* is even a possible sense of *vehicle*” (844).

One lesson to draw from Lee and Mouritsen’s results with regard to “bicycle” is that we should be quite cautious about drawing a strong conclusion concerning absence from limited

searches of a corpus. Finding that the top 50 collocates and a random sample of 100 concordance lines for “vehicle” does not indicate the use of the term to include bicycles provides reasonably compelling evidence that this is not an especially common use. But it does not provide compelling evidence that this use is *not* current or even that this use is *not* to be found in the corpus. Drawing this type of conclusion would require a more extensive analysis.

To illustrate, I will extend Lee and Mouritsen’s analysis concerning bicycles and airplanes (and to further “disputed cases” below). To do so, I’ll switch from the NOW corpus, which is comprised of texts from web-based newspapers and magazines, to COCA, which provides a more balanced corpus being drawn from wider range of sources.

A search for collocates of “vehicle” on COCA provides a similar list to that reported by Lee and Mouritsen. Here are the top 50 in order of decreasing frequency:

*motor, electric, drive, fuel, utility, driver, road, armored, st, sport, operate, steal, emergency, truck, unmanned, passenger, launch, emission, park, us, traffic, speed, hybrid, commercial, autonomous, aerial, accident, fleet, sport-utility, tank, stolen, investment, sales, strike, engine, used, all-terrain, average, recreational, ford, equipment, registration, oct, crash, license, bradley, approach, fighting, theft, off-road*⁷

As with the list produced by Lee and Mouritsen, the most frequent collocates are suggestive of the hypothesis that “vehicle” is most commonly used with regard to automobiles. And, again, we find that “airplane” and “bicycle” do not occur in the top 50 collocates. We should be hesitant to draw a strong conclusion from this, however. For one, we might also note that “automobile” does not occur in either the above list or Lee and Mouritsen’s (although “cars” did occur in their list).

Further, for purposes of assessing whether “vehicle” is ever used in the corpus to refer to such

⁷ All searches were conducted April 2021. An examination of the context for occurrences of “vehicle” and “st” reveals that this includes both “st” as an abbreviation for “Saint” (as in “all this took place a block away from Boulevard St. Laurent and the vehicle involved was a white grocery van”) and as an abbreviation for “Street” (as in “a key was used to damage a vehicle [on] King St., 4300 block”). Similarly, “us” includes both entries like “there were three of us in the vehicle” and “the company passed 200,000 in U.S. sales of electric vehicles last quarter.” Finally, context indicates that “bradley” is overwhelmingly referring to a type of military vehicle.

conveyances, as opposed to whether it is commonly used in this way, there is little reason to restrict ourselves to the top 50 collocates. In fact, expanding the list, we find that “bicycle” occurs in the top 150 collocates, appearing ahead of “automobile”; and, while “airplane” does not occur, “aircraft” is in the top 100 collocates and “helicopter” in the top 150.

Further, searching instead for collocates of “vehicle” *and* “bicycle,” including their plurals, gives 83 instances on COCA (which is comparable to “vehicle(s)” and “automobile(s)” with 55 instances and “vehicle(s)” and “airplane(s)” with 25 instances). While many of these cooccurrences do not reflect that bicycles are being described as vehicles, looking at the concordance lines we find that some of them are. For instance, we find texts such as “his bicycle is his only vehicle” and “remember that a bicycle is a legal vehicle and always ride with the traffic, not against it.” Similarly, looking at the cooccurrences of “vehicle(s)” and “airplane(s)” we find a number of texts where airplanes are being described as vehicles, including “it looks like according to the chatter, the airplane is the preferred vehicle as a weapon” and “I discovered a large, flat expanse named O’Hare Airport, to and from which great winged vehicles called airplanes somehow flew.” Thus, with regard to the question of whether “*airplane* is even a possible sense of *vehicle*,” we can offer a clear affirmative answer: “vehicle” is sometimes used in American English to refer to airplanes, even if this does not appear to be a common use.

While we do find concordance lines where “vehicle” is used to refer to a bicycle or airplane, we also find texts that suggest a contrast such as “services to reduce or eliminate your need to own your own vehicle or bicycle” or “patrol in vehicles or occasionally in airplanes.” Often, the contrast is with the phrase “motor vehicle,” such as in “open to bicycle travel and motor vehicle use.” In fact, “motor vehicle” is the most common bigram for “vehicle” and a look through concordance lines suggests that this is, in fact, what is often meant by generic uses of the term.

A potential worry, here, is that the frequency of collocates in the corpus are likely to in part reflect how frequently we talk about different types of conveyances, and this will be influenced by the prevalence of these conveyances and how newsworthy they are (Herenstein 2017).⁸ Plausibly, we might expect to see more references to automobiles than bicycles or airplanes in a standard corpus given the relative frequency of these conveyances. In fact, frequency counts from COCA suggest that this is the case, with “automobile(s)” or the more common “car(s)” occurring 324,960 times compared to 51,848 times for “bicycle(s)” or “bike(s)” and 86,794 times for “airplane(s)” or “plane(s).” The worry, of course, is that accepting that people talk about cars more frequently than bikes or planes, then even if “vehicle” could be applied equally to each, we would still expect the most frequent collocates to be suggestive of automobiles.

3. The Nonappearance Fallacy

Tobia (2020) seeks to assess the reliability of corpus methods for purposes of legal interpretation. He calls on the results of a series of questionnaire studies to conclude that “the way people understand ordinary terms and phrases (for example, ‘vehicle’ or ‘carrying a firearm’) varies systematically from what a dictionary definition or relevant legal corpus linguistics’ usage data would indicate about the meaning” (734). This leads to a striking philosophical conclusion—that “*ordinary meaning diverges from ordinary use*” (735).

At the heart of this critique is the claim that advocates of the use of corpus methods in legal interpretation, most prominently including Lee and Mouritsen, employ fallacious reasoning, generating instances of what Tobia terms *The Nonappearance Fallacy*. This is “the mistaken

⁸ The latter issue that might be compounded by the type of sources used to produce a corpus, which often include magazines and newspapers

assumption that the nonappearance of some use in a corpus indicates that this use is outside of ordinary meaning” (789). This is illustrated with an argument based on the conclusion suggested by Lee and Mouritsen above: “*across thousands of sources in our corpus, we could not find even one example of an airplane referred to as a ‘vehicle’; therefore the ordinary meaning of ‘vehicle’ does not include airplanes*” (789–790). It is important to note, here, that Lee and Mouritsen do not actually put forward such an argument; rather, they take the non-appearance of “airplane” to *raise the question* of whether this is a possible current use of “vehicle” (see also Lee and Mouritsen 2021).

Is the nonappearance fallacy actually a fallacy? A first thing to point out is that we should distinguish this “fallacy” from drawing a hasty conclusion based on a given corpus analysis. Tobia can be read as primarily critiquing how corpus methods have *typically* been used in the law, while allowing that more careful analyses might play an important role. Read in this way, Tobia’s critique is not at odds with the methodological points I am attempting to draw out here. Nonetheless, the description given for the nonappearance fallacy is suggestive of a stronger claim—one concerning what we can conclude from nonappearance in a *corpus* rather than the failure to find appearances in a limited *analysis* using that corpus.

As noted above, it is *not* the case that the use of “vehicle” to include airplanes does not occur in standard corpora. In fact, although I employed COCA, instances of such a use are also to be found in NOW, even when restricting to sources from the United States and in the timeframe employed by Lee and Mouritsen (e.g., “few commercial airplanes—the preferred vehicle for moving gold—have been flying in recent months”). The point is that it is hasty to conclude that a particular use does not occur in a corpus from an analysis based, for instance, on the top 50 collocates of the term of interest and a random sample of 100 concordance lines

(although it is fair to take such findings to raise the question). Further, we should be cautious about drawing such a conclusion from even a more detailed look at a corpus, or a range of corpora, *if* one of the terms of interest is uncommon, as will be illustrated below.

If we were to find (against the actual facts) that in large corpora like NOW and COCA there are *no* instances of a reasonably common term like “vehicle” being used to refer to a relatively common type of conveyance like an airplane, this would certainly be suggestive that the term is not currently used in this way (although it might have been in a past, which might be assessed by turning to historical corpora). And if we further found that (against the actual facts) “vehicle” was never used to refer to similar arial conveyances, such as aircraft or helicopters, this would at the very least seriously raise the question of whether such conveyances fall under the contemporary, ordinary meaning of “vehicle.” At this point, strong alternative evidence would be called for if we were to conclude that despite the lack of corpus evidence, the ordinary meaning nonetheless still includes such conveyances.

The primary evidence that Tobia provides with regard to the ordinary meaning of “vehicle” today comes from the “concept conditions” in the questionnaire studies he presents. In his first study, lay participants in the concept condition were asked to consider the noun “vehicle” and then to answer if each of ten potential types of vehicles is a vehicle, selecting either “Yes” or “No” for each. In addition to asking about “vehicle” itself, this list included both what we might consider central cases (automobile, car, bus, truck) and disputed cases (bicycle, airplane, ambulance, golf cart, and toy car). Tobia found that over 90% of participants judged that an automobile is a vehicle, and similarly for car, bus, truck, and golf cart. Roughly 85% made the same judgment about an ambulance, roughly 75% for an airplane, and roughly 65% for a bicycle. By contrast, over 75% of participants judged that a toy car is *not* a vehicle.

Similar results were found for lay participants in a subsequent study with an expanded range of potential vehicles. Participants first evaluated the 10 entities from the previous study, then a further set of 15 potential vehicles. Based on the results reported in Figure 5, for “automobile,” “car,” “bus,” “truck,” “bicycle,” “airplane,” “ambulance,” and “golf cart” a clear majority of lay participants gave an affirmative response. And a clear majority gave a negative response for “toy car.” Turning to the new items, for “moped,” “helicopter,” “horse-drawn carriage,” and “a nonfunctioning commemorative truck (e.g. a World War II Truck that has been decorated as a World War II monument)” a clear majority of lay participants gave an affirmative response. In contrast, for “a drone,” “skateboard,” “pair of roller skates,” “baby stroller,” “pair of crutches,” “pogo stick,” “baby shoulder-carrier,” “life raft,” and “zip line” a clear majority gave negative responses. Finally, participants were split with regard to an “electric wheelchair,” with a slight majority affirming, and with regard to a “wooden canoe,” with a slight majority denying.

As Lee and Mouritsen (2021) point out, Tobia treats these results as speaking directly to the ordinary meaning of “vehicle.” But, as we’ll return to below, this requires some interpretation that goes beyond the results, both with regard to how participants understood the task presented to them and how this bears on ordinary meaning. Setting this aside for the time being, however, it is worth noting that each of the potential types of vehicles that a majority of lay people affirmed as vehicles in Tobia’s studies can be rather quickly found to be attested in COCA. In fact, using the labels reported in Figure 5 (e.g., “wheelchair” instead of “electric wheelchair” and “WWII Truck” instead of the longer description), and setting aside “airplane” and “bicycle” which were discussed above, I was able to record an instance of each being described as a vehicle in COCA, with the exception of “WWII truck,” in roughly 15 minutes of total search time. Texts reflecting this usage include:

“Thursday during his arraignment on charges of repeatedly stabbing a 60-year-old tow **truck** driver whose vehicle had somehow crushed Horner’s mother”

“Petzel drove his pickup truck into the back of a **car** and sent the vehicle into the path of a motorcycle coming the other way”

“Several people were injured in a crash involving a charter **bus** and another vehicle in Baltimore County”

“Jones faces lesser criminal charges of second-degree vehicle tampering in the **automobile** theft”

“Perhaps in an **ambulance**, certainly a vehicle of some kind”

“Abrams tried to impress the fact that the **golf cart** is a vehicle and safety is of the utmost importance”

“people need to be re-educated so they look at mass transit or light vehicle (bicycles, **mopeds**, scooters or motorcycles) before they jump in their car”

“he will get into a vehicle, perhaps a **helicopter**, and he will swing by the White House”

“in loan from the school’s **carriage** shop, vehicles ranging from a phaeton to a farm wagon sat side by side”⁹

“interest by elite athletes, despite the high cost, since racing **wheelchairs** are special purpose vehicles”¹⁰

An initial search for “WWII truck” yielded no instances, although for a rather clear reason: neither “WWII truck,” “World War II truck,” or their plurals appear in COCA. This illustrates the worry raised above concerning drawing a conclusion from the nonappearance of a particular use of a word or phrase that does not occur frequently. That said, there are certainly instances of trucks being referred to as vehicles in the corpus. Further, a bit more digging reveals instances of a WWII truck being described as a vehicle, such as “the old World War II vehicle jerked

⁹ A phaeton being a type of horse-drawn carriage.

¹⁰ While, racing wheelchairs are not electric wheelchairs—the specific phrase tested by Tobia—another concordance line is at least consistent with treating electric wheelchairs as vehicles: “All feature at least one continuous mile of wide (12 to 14 feet), paved path of secure, scenic footing, where no cars or motorized vehicles are allowed (wheelchairs are OK).”

forward,” where the expanded context describes the vehicle at issue as a Jeep (for WWII this would most likely be the Willys MB / Ford GPW, which is also known as the ¼-ton 4x4 Command Reconnaissance Truck).

In contrast, I was able to find instances for only two of the remaining items from Tobia’s list being referred to as vehicles in COCA—drones (e.g., “a new brand of unmanned aerial vehicle, or drone”), as noted by Lee and Mouritsen (2018), and canoes (e.g., “our vehicle was the voyaging canoe Hokule’a”).¹¹ Thus, while there are attested instances of “vehicle” being used to refer to drones and canoes in the corpus, a majority of the lay participants in Tobia’s study denied that canoes are vehicles (roughly 55%) and a large majority denied that drones are vehicles (just over 80%). As such, while Tobia (2020: 727) holds that “ordinary meaning exceeds datasets of common usage—even very large ones,” even if we were to accept that majority opinion on the conceptual judgments he elicited is the definitive measure of ordinary meaning, it is simply not clear that this is the case.

Tobia (2020: 795) notes that “it is tempting to think that any acceptable use must be found *somewhere* in a large corpus, and any use that is not reflected is therefore not part of the ordinary meaning.” He then takes the results of his study to suggest against this tempting thought:

However, as the experimental results indicate, ordinary meaning sometimes diverges from ordinary use: people’s understanding of language is not always reflected in recorded speech and writing, especially their understanding concerning nonprototypical category membership. (790)

While it is certainly correct that we should not expect every possible use of a term to appear in even large corpora, especially if one of the target terms is infrequent, Tobia’s claim undersells

¹¹ Hōkūle’a is a replica of a traditional Polynesian voyaging canoe, and while such canoes were traditionally built with koa wood hulls, due to time-constraints the hulls of Hōkūle’a were built of plywood, fiberglass, and resin (archive.hokulea.com/build.html).

their power. In fact, each of the conveyances that his experimental results suggest as an acceptable reference for “vehicle” is attested in COCA. And this should not be overly surprising for exactly the reason Tobia gestures at in noting that this thought is tempting: we are dealing with large collections of utterances, and generally expect meaning to be reflected in use.

In fact, that the readily attested uses in COCA so closely correspond with the judgements of the majority or near-majority of Tobia’s participants offers supportive evidence for these being acceptable uses. That is, rather than finding that the corpus results and experimental results conflict with one another, they provide a consilience of evidence: all of the items that a majority of Tobia’s participants affirmed I was able to find in the corpus, and all but two of the items that a majority denied I was unable to find in the corpus. As noted above, these two items were “wooden canoe,” where participants were split, and “drone,” where a large percentage of Tobia’s participants judged are not vehicles. There are a number of possible explanations of this divergence for “drone.” It might be that we should treat the phrase “unmanned aerial vehicle”—or even uses of “vehicle” alone that are related to this phrase—as using the term in a different sense than the one of interest, such that occurrences in the corpus should be discounted. Relatedly, the divergence might reflect differences in presumed context, with “unmanned aerial vehicle” being used more commonly in a military context, including to refer to drones *transporting* munitions or cargo, while “drone” in Tobia’s studies might have instead suggested consumer drones that might be treated as more of a toy (and hence garner ratings more similar to those found for “toy car”).

This raises a larger issue with regard to the assumption that questionnaire studies like those employed by Tobia in his concept conditions can cleanly deliver the ordinary meaning of a term like “vehicle.” Lee and Mouritsen (2021) raise a number of potential concerns about this

assumption, including worries about how readily we can infer natural linguistic behaviors from a survey instrument like that employed by Tobia. A general worry here is that studies like those reported by Tobia are not free from context, but instead provide an artificial context: participants know that they are taking part in a study and are being observed, and the judgments they give can reflect this (see Sytsma and Livengood 2016, Chapter 9, for discussion). Participants' responses are elicited in an experimental context that might affect the judgments that they make.

A number of choices go into designing questionnaire studies, including the wording of the materials and the types of examples used. Focusing on the example of “drone,” we can note that the question in Tobia's studies occurred with 14 other examples and that these were given after 10 prior questions. The range of entities included in these questions might reasonably be expected to set the context for how “drone” would be understood, and it seems plausible that in a context including questions about a toy car, baby stroller, horse-drawn carriage, pogo stick, skateboard, and so on, the notion of “drone” that would occur to participants would correspond more with toys flown in a park than combat vehicles. Further, we might worry that the first 10 items leaned toward more clear examples of vehicles, with the exception of “toy car,” while the 15 items presented after this focused on less clear cases. That “drone” occurred alongside unattested examples like “baby stroller” and “pogo stick” might have encouraged skepticism with regard to whether drones are vehicles. The point here is not to discount questionnaire studies, or even the specific types of questions used by Tobia, but instead to note that they aren't free from problems.

In Tobia's second study, participants were asked to consider the noun “vehicle,” then made judgments about each of the ten entities from his first study with regard to whether it is a *prototypical* vehicle and whether it is *technically* a vehicle. While Tobia primarily discusses the

results with regard to two further conditions in his first study not discussed above, the results are also potentially telling with regard to his concept condition.¹² What Tobia finds is that the mean rating for all of the entities in the technically condition was notably above the midpoint on his scale, with the exception of “toy car,” and with “bicycle” receiving the lowest rating of the other items. This coincides with the majority responses for his concept condition, as discussed above, suggesting the possibility that Tobia’s concept condition tended to elicit judgments about what is *technically* a vehicle. In contrast, judgments of prototypicality reflected the expected central cases, with only “automobile,” “car,” “bus,” and “truck” showing mean ratings notably above the midpoint (followed by “ambulance” with a mean just above the midpoint).

With regard to “ordinary meaning” it is a matter of interpretation whether what we’re after is something more like prototypical examples of vehicles or a broader range of entities that might technically be counted as falling under the concept. The key point I want to draw out here, however, is that the data from Tobia’s second study is also consistent with the corpus analyses we have seen. The prototypical judgments coincide with the most common use of “vehicle” suggested by the analyses, while the technically judgments coincide with the less frequent but attested uses in the corpora. Again, what we find is a convergence between the questionnaire studies and the corpus analyses.

¹² In the other two conditions, participants were given either details about a dictionary definition of “vehicle” or basic corpus results (the top fifty collocates from Lee and Mouritsen along with nine concordance lines), with these being phrased in terms of a fictional type of entity—an “ailac.” Participants were then asked to make judgments about the same 10 entities as in the concept condition, but now answering, e.g., “Is a car an ailac?” Tobia takes these results to speak to the use of dictionaries and corpus linguistics in the law. As above, here it is important to distinguish between how corpus linguistics have been used in the law and how they might be used to give a more careful analysis. Tobia’s results arguably speak to the former, but not the latter. Thus, while the information Tobia provided to participants “is *precisely* what recent advocates of legal corpus linguistics recommend” (2020: 756), this information does *not* represent a careful corpus analysis of “vehicle” for the reasons discussed above, such that the relevance to assessing the *careful* use of corpus tools in legal interpretation is highly doubtful (see Lee and Mouritsen 2021 for further discussion of this point).

4. Providing Context

Klapper et al. (forthcoming) are also critical of the use of corpus methods in assessing ordinary meaning for purposes of legal interpretation, but they advocate for a different type of questionnaire method from those discussed above from Tobia.¹³ As noted, the ordinary use of a term will vary with context. And one worry about questionnaire studies is that they might inadvertently provide a context that biases participants' responses. As Lee and Mouritsen (2018: 861) state, "survey data is notoriously susceptible to context effects and response bias." One benefit of corpus methods, then, and one prominent reason that has been offered for their use in experimental philosophy, is that they provide a view of "words in the wild," outside of experimental contexts.¹⁴ This is a key part of recent claims that corpus methods can fruitfully complement the questionnaire studies run by experimental philosophers.

The flip side of this coin is that corpus methods are not so well suited to telling us about the use of a term or phrase in a given context, especially if the context is highly specific. As noted above, if the context of interest is specific enough, it is unlikely that there will be sufficient instances of the term or phrase being used in that context in standard corpora to warrant

¹³ See, however, Experiments 1A and 1B reported in the appendices to Tobia (2020). These studies included a description of an ordinance, similar to the context provided in the control condition of Klapper et al.'s study described below. In the concept condition of 1A, this was "Now imagine that a town passes an ordinance that says 'no vehicles in the park.'" Participants then rated whether each of the same ten potential vehicles used in Tobia's first study was allowed in the park. Results were similar, except that now only a slight minority treated bicycles as falling under the ordinance.

¹⁴ Often questionnaires are used to elicit conceptual judgments with regard to specific vignettes, typically drawn from thought experiments in the literature. This can then raise the question, however, of whether the results tell us something interesting about the lay concept at issue or whether the results are specific to the context provided by the vignette. For an example of relevance to the law (Tobia 2021, Knobe and Shapiro 2021), consider the large body of work indicating that norm violations impact people's causal judgments (e.g., Knobe and Fraser 2008, Hitchcock and Knobe 2009, Sytsma et al. 2012, Kominsky et al. 2015, Henne et al. 2017, Livengood et al. 2017, Kominsky and Phillips 2019, Sytsma forthcoming-b). One prominent response to this work has been to argue that this "norm effect" simply reflects experimental pragmatics, with participants inferring that the experimenters meant something like responsibility when they asked about causation (Samland and Waldmann 2016). In response, Sytsma et al. (2019) turn to corpus methods to make the case that the similarity between the use of "caused the" and "responsible for the" is found in ordinary English, outside of the experimental setting.

reasonable inferences. Questionnaire studies, on the other hand, can provide a specific context. This is of clear relevance in legal interpretation, where a statute provides a context of interest. To illustrate, we've seen above that the corpus evidence suggests that "vehicle" is used with reference to a number of different types of transportation—most commonly cars, but also other types of motor vehicles, as well as (less frequently) bicycles, helicopters and other aircraft, carriages, and so on. And this is, in fact, in line with the data Tobia reports for conceptual judgments. That the term is used to refer to this range of types of transportation, however, does not necessarily mean that people will typically interpret "vehicle" this broadly in a specific context, such as when presented with the rule "no vehicles in the park." As Klapper et al. (forthcoming: 22) note, "surveys, unlike any other methods, can provide respondents with a whole statutory provision," allowing researchers to "test a term's ordinary meaning as part of a statute and not simply isolated words."

Klapper et al. demonstrate by providing the results of a questionnaire study where participants were provided with a legal context for their conceptual judgments. In the first part of their study, participants were given a brief description of three out of five statutes involved in actual legal cases. More importantly for our purposes, in the second part participants were then given one of six vignettes describing a prohibition on vehicles in a public space. They then rated their approval of the ordinance and were asked to judge whether they thought the ordinance prohibited a series of potential vehicles: "Cars," "Bicycles," "Mopeds," "Baby carriages (strollers)," "Tobboggans," "Skateboards," "Drones," "Motorized wheelchairs," "Ambulances (even in cases of emergencies)," "Shopping carts," and "Tricycles." The vignettes varied in terms of the background motivation provided for the ordinance—including the harm it was in response to (injury, property damage, nuisance) and the conveyance involved in the harm (go-

kart or skateboard)—and the setting for the ordinance (park or college campus). A control condition simply told participants: “A local city ordinance prohibits ‘vehicles in the park.’”

Based on the corpus and questionnaire studies discussed above, what results would we expect to find here? We’ve seen that “vehicle” is most commonly used to pick out cars and other motor vehicles, and that this is in line with the empirical results for prototypicality judgments. But the term is also sometimes used more broadly to include other conveyances such as bicycles, with this being in line with the empirical results for judgments about what technically counts as a vehicle. Focusing on Klapper et al.’s control condition, the presumed setting for a prohibition against vehicles in the park would be a natural, public area focused on recreation, suggesting an area largely set apart from motor vehicle traffic.¹⁵ And the lack of further specification with regard to “vehicle,” coupled with the evidence surveyed above with regard to the most common use of the term, would reinforce this contrast. That is, it seems reasonable to expect that people would tend to focus on the most common and prototypical examples here. Based on the previous evidence, we would then expect a large majority of participants to judge that the ordinance prohibits cars and a weaker majority to judge that it prohibits mopeds (sharing features of both motorcycles and bicycles). Similar considerations would suggest some uncertainty with regard to ambulances, while predicting that the majority would deny that the ordinance applies to the remaining items. With regard to ambulances, however, the phrasing of the item deliberately emphasizes their use as emergency vehicles, and common sense would suggest that they should *not* be prohibited from responding to emergencies in a park.

These predictions are closely in line with Klapper et al.’s results. While they do not provide the percentages for each item in the control condition alone, across the conditions they

¹⁵ If there were doubts here, this could be addressed with further empirical work, including a corpus analysis focused on the noun “park.”

report finding that the top four items in decreasing percentage of agreement were cars (80%), mopeds (66%), bicycles (31%), and skateboards (29%). Further, the percentages for bicycles and skateboards, here, appear to be elevated somewhat by a notable increase in agreement for the condition in which the motivation for the prohibition involves a skateboarder crashing into a family: in this condition 60% answered that the prohibition applied to bicycles and 59% that it applied to skateboards.

As this illustrates, Klapper et al. found that the information provided about the motivation for the prohibition notably affected participants' responses. Similarly, they found that how expansive participants took the prohibition to be was negatively related to their attitudes toward the prohibition. In line with the above discussion, we might take this to highlight concerns with questionnaire methods, suggesting that participants' attitudes and the specific context provided are impacting their judgments, potentially compromising what we should derive from these results with regard to ordinary meaning.

While Klapper et al. acknowledge this worry, they argue that this does not undermine the value of questionnaire methods, noting that surveys can include evaluative questions about statutes that "can allow interpreters to isolate the effects of respondents' normative preferences... and to control for it in some degree when estimating ordinary meaning" (37). I fully agree; but, I also hold that it is important to clearly note the hedge here ("in some degree"), since it raises questions about how fully factors like this have been controlled for. Such concerns bring us back to the methodological point I've been stressing, highlighting the value of calling on distinct sources of evidence. Since the use of questionnaire methods does not preclude the use of corpus methods, and vice versa, we can further check worries about one type of method by employing another.

5. Conclusion

While recent methodological debates in experimental jurisprudence have suggested an adversarial relationship between advocates of corpus methods and advocates of questionnaire methods, the situation is rather different in experimental philosophy, where a number of practitioners have been pushing for the use of corpus methods as a complement to questionnaire methods, with the hope that by calling on distinct sources of evidence potential issues with each can be controlled for, ultimately providing a consilience of evidence. I believe this is the right approach. Methods are not Highlanders where there can be only one. Researchers can (and I think should) use both questionnaire methods and corpus methods in investigating the ordinary meaning of words or phrases of interest. This has been illustrated by taking a new look at the legal hypothetical of a prohibition on vehicles in the park. Despite three sets of authors advocating for different methods, and the suggestion that these methods are at odds, I believe that the results are instead remarkably consistent and together paint a fuller picture of the “ordinary meaning” of the term “vehicle,” both with regard to its most common or prototypical use and with regard to its broader extension or the range of conveyances that it might technically be applied to.

References

- Alfano, M. (forthcoming), “Digital humanities for history of philosophy: A case study on Nietzsche,” in L. Levenberg and T. Neilson (eds), *Handbook of Methods in the Digital Humanities*, Lanham MD: Rowman & Littlefield.
- Alfano, M. and M. Cheong (2019), “Examining Moral Emotions in Nietzsche with the Semantic Web Exploration Tool: Nietzsche,” *Journal of Nietzsche Studies*, 50(1): 1–10.
- Alfano, M., A. Higgins, and J. Levernier (2018), “Identifying Virtues and Values Through Obituary Data-Mining,” *The Journal of Value Inquiry*, 52: 59–79.

Allen, C., H. Luo, J. Murdock, J. Pu, X. Wang, Y. Zhai, and K. Zhao (2017), “Topic Modeling the Hàn diǎn Ancient Classics,” *Journal of Cultural Analytics*, 1(2).

Allen, C. and J. Murdock (forthcoming), “LDA Topic Modelling: Contexts for the History and Philosophy of Science,” in G. Ramsey and A. De Block (eds), *The Dynamics of Science: Computational Frontiers in History and Philosophy of Science*, Pittsburgh: Pittsburgh University Press.

Andow, J. (2015a), “How Distinctive is Philosophers’ Intuition Talk?” *Metaphilosophy*, 46(4-5): 515–38.

Andow, J. (2015b), “How ‘Intuition’ Exploded,” *Metaphilosophy*, 46(2): 189–212.

Betti, A., H. van den Berg, Y. Oortwijn, and C. Treijtel (2019), “History of Philosophy in Ones and Zeros,” in E. Fischer and M. Curtis (eds), *Methodological Advances in Experimental Philosophy*, 295–332, London: Bloomsbury Academic.

Bluhm, R. (2013), “Don’t Ask, Look! Linguistic Corpora as a Tool for Conceptual Analysis,” in M. Hoeltje, T. Spitzley, and W. Spohn (eds), *Was dürfen wir glauben? Was sollen wir tun? Sektionsbeiträge des achten internationalen Kongresses der Gesellschaft für Analytische Philosophie e.V.*, 7–15, DuEPublico.

Bluhm, R. (2016), “Corpus Analysis in Philosophy,” in M. Hinton (ed), *Evidence, Experiment, and Argument in Linguistics and the Philosophy of Language*, 91–109, New York: Peter Lang.

Bonino, G., P. Maffezoli, and P. Tripodi (forthcoming), “Logic in Analytic Philosophy: A Quantitative Analysis,” *Synthese*.

Byrd, R., N. Calzolari, M. Chodorow, J. Klavans, M. Neff, and O. Rizk (1987), “Tools and methods for computational lexicology,” *Computational Linguistics*, 13: 219–40.

Caton, J. (2020), “Using Linguistic Corpora as a Philosophical Tool,” *Metaphilosophy*, 51(1): 51–70.

de Villiers, J., R. Stainton, and P. Szatmari (2007), “Pragmatic Abilities in Autism Spectrum Disorder: A Case Study in Philosophy and the Empirical,” *Midwest Studies in Philosophy*, 31(1): 292–317.

Fischer, E. and P. Engelhardt (2017), “Diagnostic Experimental Philosophy,” *Teorema: International Journal of Philosophy*, 36(3): 117–37.

Fischer, E. and P. Engelhardt (2020), “Lingering stereotypes: Salience bias in philosophical argument,” *Mind and Language*, 35: 415–39.

Fischer, E., P. Engelhardt, and A. Herbelot (2015), “Intuitions and Illusions: From Explanation and Experiment to Assessment,” in E. Fischer and J. Collins (eds), *Experimental Philosophy, Rationalism, and Naturalism. Rethinking Philosophical Method*, 259–292, London: Routledge.

Fischer, E., P. Engelhardt, and J. Sytsma (2021), “Inappropriate stereotypical inferences? An adversarial collaboration in experimental ordinary language philosophy,” *Synthese*, 198(11): 10127–10168.

Fischer, E. and J. Sytsma (2021), “Zombie Intuitions,” *Cognition*, 215: 104807.

Gastaldi, J. (2021), “Why can computers understand natural language? The Structuralist Image of Language Behind Word Embeddings,” *Philosophy & Technology*, 34: 149–214.

Ginammi, A., J. Bloem, R. Koopman, S. Wang, and A. Betti (forthcoming), “Bolzano, Kant and the Traditional Theory of Concepts,” in A. de Block and G. Ramsey (eds), *The Dynamics of Science: Computational Frontiers in History and Philosophy of Science*, Pittsburgh: Pittsburgh University Press.

Hahn, U., F. Zenker, and R. Bluhm (2017), “Causal Argument,” in M. R. Waldmann (ed), *The Oxford Handbook of Causal Reasoning*, 475–94, Oxford: Oxford University Press.

Hansen, N., J. Porter, and K. Francis (forthcoming), “A Corpus Study of ‘Know’: On the Verification of Philosophers' Frequency Claims about Language,” *Episteme*.

Hart, H. L. A. (1958), “Positivism and the Separation of Law and Morals,” *Harvard Law Review*, 71(4): 593–629.

Henne, P., Á. Pinillos, and F. De Brigard (2017), “Cause by Omission and Norm: Not Watering Plants,” *Australasian Journal of Philosophy*, 95(2): 270–83.

Herbelot, A., E. von Redecker, and J. Müller (2012), “Distributional techniques for philosophical enquiry,” *Proceedings of the 6th Workshop on Language Technology for Cultural Heritage, Social Sciences, and Humanities*, 45–54.

Herenstein, E. (2017), “The Faulty Frequency Hypothesis: Difficulties in Operationalizing Ordinary Meaning Through Corpus Linguistics,” *Stanford Law Review Online*, 70: 112–122.

Hinton, M. (2020), “Corpus Linguistics Methods in the Study of (Meta)Argumentation,” *Argumentation*.

Klapper, S., S. Schmidt, and T. Tarantola (forthcoming), “Ordinary Meaning from Ordinary People,” *UC Irvine Law Review*.

Knobe, J. and B. Fraser (2008), “Causal judgments and moral judgment: Two experiments,” in W. Sinnott-Armstrong (ed), *Moral Psychology, Volume 2: The Cognitive Science of Morality*, 441–47, Cambridge: MIT Press.

- Knobe, J. and J. Prinz (2008), “Intuitions about consciousness: Experimental studies,” *Phenomenology and the Cognitive Sciences*, 7: 67–85.
- Knobe, J. and S. Shapiro (2021), “Proximate Cause Explained: An Essay in Experimental Jurisprudence,” *University of Chicago Law Review*, 88: 165–236.
- Kominsky, J. and J. Phillips (2019), “Immoral Professors and Malfunctioning Tools: Counterfactual Relevance Accounts Explain the Effect of Norm Violations on Causal Selection,” *Cognitive Science*, 43(11): e12792.
- Kominsky, J., J. Phillips, T. Gerstenberg, D. Lagnado, and J. Knobe (2015), “Causal superseding,” *Cognition*, 137: 196–209.
- Lean, O., L. Rivelli, and C. Pence (forthcoming), “Digital Literature Analysis for Empirical Philosophy of Science,” *British Journal for the Philosophy of Science*.
- Lee, T. and S. Mouritsen (2018), “Judging Ordinary Meaning,” *The Yale Law Journal*, 127(4): 788–879.
- Lee, T. and S. Mouritsen (2021), “The Corpus and the Critics,” *The University of Chicago Law Review*, 88: 275–366.
- Levinson, S. (2000). *Presumptive Meanings. The Theory of Generalized Conversational Implicature*, Cambridge, Mass.: MIT Press.
- Liao, S., L. McNally, and A. Meskin (2016). “Aesthetic Adjectives Lack Uniform Behavior,” *Inquiry: An Interdisciplinary Journal of Philosophy*, 59(6): 618–31.
- Livengood, J., J. Sytsma, and D. Rose (2017), “Following the FAD: Folk attributions and theories of actual causation,” *Review of Philosophy and Psychology*, 8(2): 274–94.
- Ludlow, P. (2005), “Contextualism and the New Linguistic Turn in Epistemology,” in G. Preyer and G. Peter (eds), *Contextualism in Philosophy: Knowledge, Meaning, and Truth*, 11–51, Oxford: Oxford University Press.
- Malaterre, C., J. Chartier, and D. Pulizzotto (2019), “What is this thing called *Philosophy of Science*? A computational topic-modeling perspective, 1934–2015,” *HOPOS: The Journal of the International Society for the History of Philosophy of Science*, 9(2): 215–49.
- Malaterre, C. and J. Chartier (forthcoming), “Beyond Categorical Definitions of Life: A Data-Driven Approach to Assessing Liveness,” *Synthese*.
- Malaterre, C., F. Lareau, D. Pulizzotto, and J. St-Onge (forthcoming), “Eight journals over eight decades: a computational topic-modeling approach to contemporary philosophy of science,” *Synthese*.

- Malaterre, C., D. Pulizzotto, and F. Lareau (2020), “Revisiting Three Decades of Biology and Philosophy: A Computational Topic-Modeling Perspective,” *Biology & Philosophy*, 35(5).
- McKinnon, A. (1970), *The Kierkegaard Indices*, Leiden: Brill.
- McRae, K. and M. Jones (2013), “Semantic memory,” in D. Reisberg (ed), *Oxford Handbook of Cognitive Psychology*, Oxford: Oxford University Press.
- Mejía-Ramos, J., L. Alcock, K. Lew, P. Rago, C. Sangwin, and M. Inglis (2019), “Using Corpus Linguistics to Investigate Mathematical Explanation,” in E. Fischer and M. Curtis (eds), *Methodological Advances in Experimental Philosophy*, 239–63, London: Bloomsbury Academic.
- Meunier, J., D. Forest, and I. Biskri (2005), “Classification and Categorization in computer-assisted reading and text analysis,” in H. Cohen and C. Lefebvre (eds), *Handbook of Categorization in Cognitive Science*, 955–78, Amsterdam: Elsevier.
- Meunier, J., S. Rolland, and F. Daoust (1976), “A system for text and content analysis.” *Computers and the Humanities*, 10(5): 281–86.
- Mizrahi, M. (forthcoming-a), “Conceptions of scientific progress in scientific practice: An empirical study,” *Synthese*.
- Mizrahi, M. (forthcoming-b), “Theoretical Virtues in Scientific Practice: An Empirical Study,” *British Journal for the Philosophy of Science*.
- Mizraha, M. (2020a), “The Case Study Method in Philosophy of Science: An Empirical Study,” *Perspectives on Science*, 28(1): 63–88.
- Mizrahi, M. (2020b), “Hypothesis Testing in Scientific Practice: An Empirical Study,” *International Studies in the Philosophy of Science*, 33(1): 1–21.
- Mizrahi, M. and M. Dickinson (forthcoming), “The Analytic – Continental Divide in Philosophical Practice: An Empirical Study,” *Metaphilosophy*.
- Murdock, J., C. Allen, and S. DeDeo (2017), “Exploration and Exploitation of Victorian Science in Darwin’s Reading Notebooks,” *Cognition*, 159: 117–26.
- Nagel, J. (2013), “Knowledge as a Mental State,” *Oxford Studies in Epistemology*, 4: 275–310.
- Nagel, J. (2021), “The Psychological Dimension of the Lottery Paradox,” in I. Douven (ed), *Lotteries, Knowledge, and Rational Belief: Essays on the Lottery Paradox*, 48–73, Cambridge: Cambridge University Press.
- Nichols, S., S. Kuman, T. Lopez, A. Ayars, and H. Chan (2016), “Rational Learners and Moral Rules,” *Mind & Language*, 31(5): 530–54.

Nichols, S. and Á. Pinillos (2018), “Skepticism and the Acquisition of ‘Knowledge’,” *Mind & Language*, 33(4): 397–414.

Nichols, S. (2021), *Rational Rules: Towards a Theory of Moral Learning*, Oxford: Oxford University Press.

Nie, Chenwei (forthcoming), “Can a bodily theorist of pain speak Mandarin?” *Philosophia*.

Overton, J. (2013), “‘Explain’ in Scientific Discourse,” *Synthese*, 190(8): 1383–405.

Pease, A., A. Aberdein, and U. Martin (2019), “Explanation in Mathematical Conversations: An Empirical Investigation,” *Philosophical Transactions of the Royal Society A*, 377: 20180159.

Pence, C. (2016), “RLetters: A Web-Based Application for Text Analysis of Journal Articles,” *PLoS ONE*, 11(1): e0146004.

Pence, C. (forthcoming), “How Not to Fight About Theory: The Debate Between Biometry and Mendelism in Nature, 1890-1915,” in A. de Block and G. Ramsey (eds.), *The Dynamics of Science: Computational Frontiers in History and Philosophy of Science*, Pittsburgh: Pittsburgh University Press.

Pence, C. and G. Ramsey (2018), “How to Do Digital Philosophy of Science,” *Philosophy of Science*, 85(5): 930–41.

Ramsey, G. and C. Pence (2016), “evoText: A New Tool for Analyzing the Biological Sciences,” *Studies in History and Philosophy of Science Part C*, 57: 83–7.

Reuter, K. (2011), “Distinguishing the appearance from the reality of pain,” *Journal of Consciousness Studies*, 18(9-10): 94–109.

Sainte-Marie, M., J. Meunier, N. Payette, and J. Chartier (2011), “The concept of evolution in the *Origin of Species*: a computer-assisted analysis,” *Literary and Linguistic Computing*, 26(3): 329–34.

Samland, J. and M. Waldmann (2016), “How prescriptive norms influence causal inferences,” *Cognition*, 156: 164–76.

Schwitzgebel, E. and C. Dacey Jennings (2017), “Women in Philosophy: Quantitative Analyses of Specialization, Prevalence, Visibility, and Generational Change,” *Public Affairs Quarterly*, 31(2): 83–105.

Slingerland, E. and M. Chudek (2011), “The Prevalence of Mind–Body Dualism in Early China,” *Cognitive Science*, 35(5): 997–1007.

Sytsma, J. (forthcoming-a), “Crossed Wires: Blaming Artifacts for Bad Outcomes,” *The Journal of Philosophy*.

Sytsma, J. (forthcoming-b), “Causation, Responsibility, and Typicality,” *Review of Philosophy and Psychology*.

Sytsma, J., R. Bluhm, P. Willemsen, and K. Reuter (2019), “Causal Attributions and Corpus Analysis,” in E. Fischer and M. Curtis (eds), *Methodological Advances in Experimental Philosophy*, 209–238, London: Bloomsbury Academic.

Sytsma, J. and J. Livengood (2016), *The Theory and Practice of Experimental Philosophy*. Peterborough: Broadview Press.

Sytsma, J., J. Livengood, and D. Rose (2012), “Two types of typicality: Rethinking the role of statistical typicality in ordinary causal attribution,” *Studies in History and Philosophy of Biological and Biomedical Sciences*, 43: 814–20.

Sytsma, J. and K. Reuter (2017), “Experimental Philosophy of Pain,” *Journal of Indian Council of Philosophical Research*, 34(3): 611–28.

Sytsma, J. and M. Snater (forthcoming), “Consciousness, Phenomenal Consciousness, and Free Will,” in P. Henne and S. Murray (eds), *Advances in Experimental Philosophy of Action*, London: Bloomsbury Academic.

Tallant, J. (2013), “Intuitions in Physics,” *Synthese*, 190: 2959–80.

Transwell, F. and M. Inglis (forthcoming). “The Language of Proofs: A Philosophical Corpus Linguistics Study of Instructions and Imperatives in Mathematical Texts,” in B. Sriraman (ed), *Handbook of the History and Philosophy of Mathematical Practice*, Springer.

Tobia, K. (2020), “Testing Ordinary Meaning,” *Harvard Law Review*, 134(2): 726–806.

Tobia, K. (2021), “Law and the Cognitive Science of Ordinary Concepts,” in B. Brożek, J. Hage, and N. Vincent (eds), *Law and Mind: A Survey of Law and the Cognitive Sciences*, 86–96, Cambridge: Cambridge University Press.

Tsugita, S., Y. Izumi, and M. Mizumoto (forthcoming), “Knowledge-How Attribution in English and Japanese,” in *Epistemic Agency and Epistemic Environments in East-West Philosophy—Extending Knowledge*, Palgrave.

Ulatowski, J., D. Weijers, and J. Sytsma (2020), “Corpus Methods in Philosophy,” *The Brains Blog*: <https://philosophyofbrains.com/2020/12/15/cognitive-science-ofphilosophy-symposium-corpus-analysis.aspx>

van Wierst, P., S. Hofstede, Y. Oortwijn, T. Castermans, R. Koopman, S. Wang, M. Westenberg, and A. Betti (2018), “BolVis: Visualization for Text-based Research in Philosophy,” *Proceedings of the 3rd Workshop on Visualization for the Digital Humanities*.

Vetter, B. (2014), “Dispositions without Conditions,” *Mind*, 123: 129–56.

Weatherson, B. (2020), *A History of Philosophy Journals, Volume 1: Evidence from Topic Modeling, 1876-2013*: <http://www-personal.umich.edu/~weath/lda/>

Willemsen, P., L. Baumgartner, S. Frohofer, and K. Reuter (forthcoming), “Examining Evaluativity in Legal Discourse: A Comparative Corpus-Linguistic Study of Thick Concepts,” S. Magen and K. Prochownik (eds), *Advances in Experimental Philosophy of Law*.

Wright, J. and K. Bartsch (2008), “Portraits of Early Moral Sensibility in Two Children’s Everyday Conversations,” *Merrill-Palmer Quarterly*, 54(1): 56–85.

Wright, J., T. Sedlock, J. West, K. Saulpaugh, and M. Hopkins (2016), “Located in the think of it: Young children’s use of thin moral concepts,” *Journal of Moral Education*, 45: 308–323.