Are moral predicates subjective? A corpus study

Isidora Stojanovic (Institut Jean Nicod, CNRS, DEC, ENS, PSL)
and Louise McNally (U. Pompeu Fabra)[1]

Forthcoming in Bordonaba, David (Ed.) *Experimental Philosophy of Language: Perspectives, Methods and Prospects*, Springer

Abstract

The nature of moral judgments, and, more specifically, the question of how they relate, on the one hand, to objective reality and, on the other, to subjective experience, are issues that have been central to metaethics from its very beginnings. While these complex and challenging issues have been debated by analytic philosophers for over a century, it is only relatively recently that more interdisciplinary and empirically-oriented approaches to such issues have begun to see light. The present chapter aims to make a contribution of that kind. We will present the results of an empirical – specifically, corpus linguistic – study that offers evidence that moral predicates exhibit hallmarks of subjectivity at the linguistic level, but also, that they differ significantly from paradigmatically subjective predicates.

1. Introduction

The nature of moral judgments, and, more specifically, the question of how they relate, on the one hand, to objective reality and, on the other, to subjective experience, are issues that have been central to metaethics from its very beginnings. While these complex and challenging issues have been debated by analytic philosophers for over a century (and by philosophers *tout court* since Plato and Aristotle), it is only relatively recently that more interdisciplinary and empirically-oriented approaches to such issues have begun to see light. The present chapter aims to make a contribution of that kind.

We begin in section 2 by contextualizing the study in the literature on moral predicates and subjectivity. We briefly look at recent studies from experimental moral philosophy that suggest that moral judgments are more subjective than factual judgments, but less so than judgments on matters of aesthetic preference and personal taste. We also look at how subjectivity has been approached in semantics and, in particular, at the idea that embedding a predicate under subjective attitude verbs like "find" can serve as a criterion for subjectivity.

Section 3 presents the corpus study. In a nutshell, we have looked at how basic moral predicates - "moral" and "immoral", "ethical" and "unethical" - as well as predicates modified by "morally" and "ethically", behave with respect to the verbs "find" and "consider", both of which denote subjective attitudes, but of different kinds. Section 4 discusses the theoretical implications of the results of our study, and argues that moral predicates exhibit hallmarks of subjectivity at the linguistic level, but are also importantly different from paradigmatically subjective predicates such as "fun" and "boring" or "delicious" and "disgusting". While the latter are clearly associated with "find"-like attitudes, the former show a preference for "consider"-like attitudes.

[1] The authors have contributed equally to the paper.

2. Moral predicates and subjectivity: A snapshot of a long-standing philosophical debate and the more recent empirical turn

2.1. The vexed issue of moral subjectivity

Moral realism, that is, the view that "moral claims [...] purport to report facts and are true if they get the facts right" (Sayre-McCord 2005, p. 1), is a long-standing position in metaethics, but also one that faces many challenges (see, i.a. Sayre-McCord 2005; Railton 2017). Consider the following two claims:

(1) a. Mao Zedong's Cultural Revolution was unethical.
b. Mao Zedong's Cultural Revolution took place in China from 1966 to 1976.

In its somewhat caricatural form, a moral realist holds that (1a) and (1b) are completely on a par: both purport to refer to objective facts, and whether they are true or false can be determined simply by looking at how things are, or were, in reality. However, while (1b) indeed reports a historical fact that can be easily confirmed as true, it is far from obvious what kind of fact would play the same role for (1a). Relatedly, disagreements on the two kinds of claims differ importantly. When and where the Cultural Revolution took place is hardly open to disagreement, and if it were, historians would be joining their forces in order to find the best way of establishing when and where it took place. Moral disagreements, in contrast, often (though not always) tend toward the unresolvable, and moreover, are arguably such because of their very nature. A hard-core communist in China in the late 1960s who supported the Cultural Revolution and a hard-core opponent of communism differ precisely in that they endorsed radically opposed systems of values, and are unlikely to be ever able to reach an agreement over a statement such as (1a).

From the basic observation that moral claims like (1a) and factual claims like (1b) do not appear to be on a par – at least, not when taken at face value – one sees a proliferation both of alternatives to and refinements of moral realism (for overview, see e.g. Soria Ruiz, Cepollaro and Stojanovic 2021). Among the alternatives, expressivist views (see e.g. Camp 2017 for overview) hold that the function of moral statements is not to express factual information, but rather to convey subjective attitudes with respect to a moral issue; for example, (1a) expresses the speaker's disapproval of the Cultural Revolution. But expressivism also faces serious challenges. For one, subjective attitudes such as (dis)approval do not bear truth value, while moral claims, at least prima facie, do. One who is in disagreement with (1a) can reply "That's not true!" Similarly, claims such as (1a) easily appear in contexts where they must be able to bear a truth value, as in (2):

(2) If the Cultural Revolution was unethical, then its leaders do not deserve monuments in their honor.

The challenge of explaining how a sentence such as (1) can express an attitude such as disapproval when uttered on its own, yet occur in complex constructions such as (2), is an instance of the so-called "Frege-Geach problem" and has been widely discussed in

metaethics (see e.g. Woods 2017 for overview). In particular, a family of accounts known as *hybrid* expressivist accounts aim to preserve the basic insights of expressivism while being able to account for the compositional behavior of moral language (see e.g. Björkholm 2022 for overview and discussion).

To bring the point home, moral statements appear to be neither (completely) objective nor (completely) subjective. This in-between status of moral statements is further confirmed by several recent empirical studies. The most famous are by Goodwin and Darley (2008, 2010, 2012). In their pioneering 2008 study, they investigated folk judgments regarding moral statements and made a four-tiered comparison between factual statements, statements on matters of social convention, statements on matters of personal taste and aesthetics, and moral statements. The task with which they presented their participants consisted of three steps. First, participants were presented with a statement and asked to which degree they agreed with the statement. Next, participants had to decide whether they thought it was a true statement, a false statement, or a matter of opinion. Finally, after a task of distraction during which the examiner would select five statements with which the participant agreed to a very high degree, participants were told that somebody else disagreed with these statements and were asked whether they thought the other person was surely wrong, whether it was possible that neither of them was wrong, or whether it was possible that the participant was wrong and the other person right. Combining these different measures, Goodwin and Darley created “a scale of objectivity” and found that moral statements were judged to be less objective than factual statements, but more objective than statements on matters of social convention or personal taste (with the former being judged more objective than the latter). In follow-up studies, they further showed that the degree of perceived subjectivity varied significantly among moral judgments themselves. Thus judgments regarding what they call contested value of life issues, such as the permissibility of abortion or euthanasia, were found to be more subjective than others.

These initial studies set into motion a rich and still incredibly active research agenda, whose goal is twofold. First, it applies experimental methodology to the study of folk judgments regarding the subjective vs. objective character of moral statements. Second, it aims to examine the implications of these empirical findings for theoretical issues discussed in metaethics. For further studies and discussion, see i.a. Wright, Grandjean and McWhite (2013), Beebe and Sackris (2016), Pölzler (2017), Pölzler and Wright (2020a, 2020b), and Sarkissian (2016). While space precludes summarizing these latter works here, the question of whether, and to what extent, moral judgments and moral statements are objective rather than subjective is still very much a matter of controversy. The study we present in Section 3 is a contribution to this debate.

2.2. Tracking subjectivity semantically

While experimental moral philosophy and psychology were concerned with subjectivity specifically with respect to morality, other issues concerning subjectivity emerged independently and came to be topics of great interest in semantics. More precisely, for the past two decades, there has been a growing interest in the semantics of subjective and evaluative predicates, among which the predicates of personal taste (PPTs, for short), such

as “tasty”, “delicious”, “disgusting”, “fun” and “boring”, have been in the center of attention; see e.g. Lasersohn (2005), Stephenson (2007), Stojanovic (2007) for early discussions, and Umbach (2021), Stojanovic and Kaiser (2022: section 2) and Willer (to appear) for recent overviews. One of the things that sparked such a vivid interest in subjective predicates is the idea of faultless disagreement; that is, the idea that it makes sense to disagree over matters of personal taste. For example, we may debate whether Monopoly is fun or boring, even if we know that Monopoly can be fun to some people, or on some occasions, and boring to other people, or on other occasions. Faultless disagreement has generated an impressive amount of philosophical literature, and continues to be a hotly debated issue; for overviews, see, i.a., Bordonaba (2017), Stojanovic (2017) Karczewska (2019), Zakkou (2019), or Zeman (2020).

While PPTs are particularly prone to generating scenarios of faultless disagreement, it has been widely noted in the literature that the phenomenon appears to be much broader than matters of personal taste. For example, vague predicates (that is, those that generate instances of the sorites paradox) and relative gradable predicates in general, when used in a positive (as opposed to comparative or superlative) form, can easily give rise to what appears to be a faultless disagreement. Consider two friends who disagree over whether a 10€ bottle of wine is expensive. One of them can judge it to be expensive because they are used to buying wine that would cost 4€ a bottle, and the other, not expensive because their standard of reference is wines that cost over 15€ a bottle. Any predicate whose application makes reference to standards that can vary across contexts can, in principle, give rise to dialogues that take the form of a faultless disagreement (see e.g. Kennedy 2013, Solt 2018, Odrowąż-Sypniewska 2021; see also Verheyen, Dewill and Egré 2018 for an experimental study of the subjectivity in gradable adjectives). What is more, even expressions such as *athlete* or *publication* can arguably generate faultless disagreement (see Sundell 2011, Stojanovic 2012) because the conditions for their application are not firmly settled, so that competent language users may disagree whether someone counts as an athlete or whether something counts as a publication without there being a completely objective way to settle the matter.

Importantly for the present purposes, when it comes to moral predicates, there is considerable controversy as to whether they are prone to faultless disagreement or not. We have already seen that empirical studies such as Goodwin and Darley (2008, 2012) provide a mitigated answer: moral statements are more prone to faultless disagreement than factual statements, but significantly less so that statements on matters of personal taste, and among moral statements themselves, some are more prone than others. Similar findings are reported in the recent study in Soria Ruiz and Faroldi (2020), while Stojanovic (2019) argues, on more theoretical grounds, that disagreements over moral issues pattern differently from disagreement over matters of taste.

Because of the elusive character of faultless disagreement, scholars have looked for other ways of identifying subjectivity at the semantic level. One diagnostic that has gained great popularity is the so-called “find” test (see Sæbø 2009). The basic idea is that it is fine to embed subjective predicates under verbs of subjective attitude such as the English "find",

but not so with nonsubjective predicates, as illustrated by the following contrast (from Kennedy 2013, p. 260):

(3) a. Anna finds trippa alla romana tasty.
b. ??Anna finds trippa alla romana to be vegetarian.

While embeddability under subjective attitude verbs such as "find" is often used as a test in order to identify PPTs, the test is not without problems. A particularly pressing problem is that there are predicates that seem to fall into something of a gray zone: while they are not outright infelicitous under "find", they are not perfectly felicitous either. Thus McNally and Stojanovic (2017) write:

> Another sign that "find" anti-selects for strictly evaluative predicates is the oddness of assertions like (4), in comparison to the more natural embedding under "consider" in (5).
>
> (4) a. ?I find Miró's mosaic on the Rambles mediocre.
> b. ?I find lying bad/worse than stealing.
> (5) a. I consider Miró's mosaic on the Rambles mediocre.
> b. I consider lying bad/worse than stealing.
>
> Though (4b) is not unacceptable, it strongly implies that the speaker has made his or her evaluation about lying on the basis of specific experiences of doing it. (2017, p. 29; numbering of examples adjusted)

Moral adjectives are among the adjectives in the gray zone. Some authors, such as Franzén (2020) and Silk (2021), take moral adjectives to be felicitous under "find", and take this as evidence to the effect that moral adjectives are subjective, while other authors, such as McNally and Stojanovic (2017) and Stojanovic (2019) take them to be marked under "find", or as noted in the cited passage, felicitous only in the context of specific subjective experiences that are compatible with, but not inherent to, moral judgments. They thus take the "find" test to offer evidence that there is an important semantic difference between moral adjectives and PPTs. This continuing controversy surrounding the embeddability of moral adjectives under "find" and its implications is thus one of the motivations for the corpus study that we have conducted and present in the next section.

Finally, even if moral predicates do not pattern in quite the same way as PPTs when it comes to faultless disagreement and embeddability under "find", this does not mean that these predicates are completely objective either. The empirical studies of Goodwin and Darley (2008, 2012) and Soria Ruiz and Faroldi (2020) show that moral predicates still elicit a significantly high intuition of faultlessness. Furthermore, moral predicates are clearly felicitous under other verbs of subjective attitude, and in particular, the verb "consider", which anti-selects for fully objective predicates (see Lasersohn 2009, Fleisher 2013, Kennedy and Willer, forthcoming):

(6) #Anna considers the sum of two and two greater than four.

Let us take stock. While it was PPTs that triggered the interest in subjectivity from a semantic point of view, one challenge that followed immediately was to know how far this notion of subjectivity extended. Appeal to the idea of faultless disagreement suggested that there were many more expressions beyond PPTs that were subjective, but the "find" test narrowed down again the range of putative subjective expressions. However, a major issue was, and continues to be, that for many expressions, the applicability of that test yields controversial results. To date, the data discussed on moral predicates and "find" have been, to our knowledge, entirely anecdotal and constructed for the purposes of making an argument. While such data can be legitimately used, the controversies described above suggest that it could be informative to take a broader, more systematic look at naturally occurring examples.

3. The corpus study

3.1. Corpus used and raw data collection method

We took a snapshot of the distribution of moral adjectives with "find" through a study carried out on the Corpus of Contemporary American English (COCA), using the search tool at www.english-corpora.org (Davies 2008). COCA has over 1 billion words spanning the years 1990-2019 and offers a sample of English evenly distributed across language drawn from academic journals, magazines, newspapers, fiction, spoken language (TV and radio interviews), TV and movie subtitles, blogs and other web pages. In this respect, it constitutes what corpus linguists would consider a balanced, representative sample of the language.

Since this is, to our knowledge, the first corpus study on this topic,[2] and given the large size and the somewhat limited linguistic information that can be searched for in the corpus using the web search tool, we opted for a limited study, in the hope of inspiring future research on broader sets of data. We focused on uses of adjectives expressing moral judgments as predicative complements to "find" and, for comparison, "consider". We chose the adjectives "moral", "immoral", "ethical" and "unethical", which we considered to be simultaneously among the most prototypical examples of adjectives used for moral judgments and the least polysemous. Of course, other adjectives can be used to make moral judgments – candidates we considered include "good", "bad", "right", "wrong", and some additional, more specific adjectives like "(un)acceptable". However, after some initial searches, we found that these all raised concerns due to their polysemy: Something can be good, bad, or (un)acceptable for moral or ethical reasons, or for other reasons not related to moral judgments. We wished to avoid having to make qualitative decisions concerning the interpretation of such adjectives. As an alternative, to broaden our dataset somewhat, we added to the search complements of the form ["morally"/"ethically" ADJECTIVE], assuming that a speaker who

[2] While philosophers in metaethics and philosophy of language are increasingly eager to look at empirical evidence concerning morality and moral language, the main focus has been on collecting data through controlled experiments (e.g. eliciting acceptability judgments), rather than from corpora. A notable exception is the corpus study presented in Reuter et al. (ms.), who use corpus data to argue that thick and thin evaluative (specifically moral) expressions are distinguishable from other types of expressions in terms of how they combine with intensifiers ("truly", "really", "very").

chooses to use the qualifiers “morally” or “ethically” is making explicit the nature of their judgment.

All words in the online version of COCA are tagged for lemma (that is, the basic form that covers all inflected forms, such as “finds”, “found”, and “finding” for “find”), as well as for part of speech (noun, verb, etc.). However, COCA is not syntactically parsed: there is no way, for example, to distinguish adjectives used as predicates (as in “We found that **unethical”**) from those used as modifiers of nouns (as in “We found that **unethical** politician collecting bribes”). We therefore could not search for adjectives specifically used as predicative complements to “find” or “consider”. The practical alternative that allowed us to collect the most examples was a search for the lemma for each verb within 9 words – the maximum window afforded by the search tool – to the left of each adjective and adverb.[3] This strategy guaranteed that we could capture examples in which the complement to the verb to which the moral adjective or adverb was ascribed was quite long (e.g. “It is difficult to consider **the employee of a company** immoral”) or where adverbial or other material intervened (e.g. “those who consider it **entirely** immoral”). However, it also meant that we collected a lot of false positives which had to be filtered out (for example, “[Y]ou don’t say whether you **consider** eavesdropping to be a moral or **ethical** act”, where “ethical” modifies “act”, and moreover where the judgment is not about whether eavesdropping is ethical or not, but rather whether it is an act of an ethical nature). We offer additional details on the data filtering in the following subsection.

Of course, our results will be better interpretable if we also have information about how adjectives behave with “find” and “consider” more generally. This requires having a sense of a) the base frequency of the two verbs and the moral adjectives (particularly when used as predicates); and b) the range and frequency of the other adjectives that occur as predicative complements to “find” and “consider”, as well as the base frequency of these latter adjectives, again, particularly as predicates. As already noted, because the corpus is not syntactically parsed, it is not trivial to extract this information reliably. However, we did attempt a broader quantitative comparison in two ways.

First, we carried out an additional search to get a sense of how the frequencies of “(im)moral” and “(un)ethical” with “find” and “consider” compare with those of other adjectives that serve as complements to these verbs. To keep the data collection manageable and as comparable as possible, we collected frequency counts for all adjectives that occurred in the context of the lemmas for “find” or “consider” followed directly by the pronoun “it”, specifying in addition that the expression directly following the adjective not be a noun, to avoid picking up uses of the adjective as a modifier. As “it” is unambiguously a pronoun (unlike “that”, which also has a use as a determiner, as in “**that** ethical dilemma”), we minimized the collection of irrelevant examples – any adjective following “it” and not followed by a noun is highly likely to be a predicate, as in “consider it ethical”. At the same time, “it” is a highly frequent word, and therefore considered likely to produce a sufficient number of hits to allow for some preliminary analysis.

---

[3] This sort of search is carried out using the collocation search option in the tool at www.english-corpora.org. For technical reasons, it was not possible to use an equivalent and intuitively more natural strategy of searching for the adjective within the same window to the right of the verb lemma.

Second, we examined the automatically calculated mutual information (MI) scores available in COCA for the verbs “find” and “consider” with all adjectives, as well as the mutual information scores for “moral”, “immoral”, “ethical”, “unethical”, “morally”, and “ethically” with all verbs. We provide further details on MI and why we looked at it in the next section.

3.2. Initial results and data filtering

The raw numbers of hits produced by first searches specifically for “(im)moral”, “(un)ethical”, “morally”, and “ethically”, are summarized in Table 1.

	moral	immoral	ethical	unethical	morally ADJ	ethically ADJ
FIND	376	64	156	56	176	32
CONSIDER	334	145	177	74	100	25

Table 1: Raw occurrences in COCA of the lemmas for “find” and “consider” within a 9-word window to the left of “(im)moral”, “(un)ethical”, “morally”, and “ethically”.

However, these had to be filtered to eliminate duplicates as well as to restrict results to uses of the adjectives as complements to the verbs in question, and of the adverbs as modifiers of adjectival complements to the verbs. This required reading the examples individually, and was carried out by McNally, a trained linguist and native speaker of English (although the task did not present any particular difficulty). The output of this process includes all examples in which the syntactic functions of the expressions are respected, even if the surface word order varies (e.g. “consensual behavior he considers immoral”, where “immoral” is ascribed to “consensual behavior” from within a relative clause, or “is considered **wrong ethically”**, a marked but grammatical option in English). Examples in which the moral adjective complement was preceded by as, a stylistic option in English, were also left in (e.g. “Neocons **consider** lying as a standard operating procedure **as perfectly ethical”**), as were those in which the moral adjective appeared as the complement to “to be” in an infinitival complement to the verb (e.g. “We **consider** her actions **to be immoral”**). In the latter case, though the syntactic structure is technically different, we see no nuance of semantic or pragmatic difference of any sort. The results, after this initial filtering, appear in Table 2.

	moral	immoral	ethical	unethical	morally ADJ	ethically ADJ
FIND	4	45	11	25	138	15
CONSIDER	32	125	32	64	70	17

Table 2: Number of occurrences in COCA of the lemmas for “find” and “consider” within a 9-word window to the left of “(im)moral”, “(un)ethical”, “morally, and “ethically”, after filtering.

As can be seen, filtering considerably reduced the number of examples. We provide further, qualitative commentary on these in the next section.

As noted at the end of the last section, in order to put these results in context, it is relevant to take into account the overall frequency of both the two verbs and the individual adjectives and adverbs. As a first approximation, we carried out a further search for the lemmas for "find"/"consider", followed immediately by "it", then directly by any word of the category adjective, and then any category other than a noun (the specific search strings used were "FIND it ADJ -NOUN" and "CONSIDER it ADJ -NOUN", where "-" is a Boolean negation operator). Due to imprecisions in the tagging, these searches also yielded various false positives that had to be manually filtered. These fell into two cases: a) examples where a noun appeared after the adjective (e.g. "consider it real **progress"**); and b) where the third item was not an adjective (e.g. "finding it **–** while", where the dash punctuation constitutes the third item). As before, this filtering was carried out by McNally. It is worth pointing out that this filtering, in the case of "find", leaves in examples that probably correspond to a distinct sense of the verb that does not involve subjective judgment, namely examples like (7).

(7) He opened the lid and **found it empty.** The two remaining seeds were gone.

In this example, the verb describes not a judgment but an event of encountering something in an objective state. This sense of "find" is highly salient with "empty" and a few other adjectives, such as "full", "intact", "(un)occupied", and "vacant". However, it cannot be reliably identified solely by considering the adjective alone – for example, one could use (8) to express a subjective judgment about a theater after a very poorly attended performance.

(8) I found it empty.

Since the examples we extracted were too numerous to verify individually, in some cases it is difficult to determine the interpretation of the verb with certainty, and overall the adjectives we considered likely to yield this interpretation constituted not more than 100 tokens, or about 0,7% of the total for "find", we chose not to exclude them.

After filtering, a total of 14536 tokens of FIND "it" ADJ and 831 tokens of CONSIDER "it" ADJ remained. We also searched for instances of "morally" or "ethically" ADJ in the same context. The results, including the number of tokens involving the four moral adjectives of interest are summarized in Table 3.

	ADJ	moral	immoral	ethical	unethical	morally ADJ	ethically ADJ
FIND it _	14536	1	6	6	2	12	2
CONSIDER it _	831	1	7	2	6	5	1

Table 3. Number of occurrences in COCA 1) of all adjectives (including, "(im)moral" and "(un)ethical"), 2) of "(im)moral" and "(un)ethical", and 3) of "morally"/"ethically" ADJ directly following FIND/CONSIDER "it" and not followed by a noun, after filtering.

For comparison we extracted two additional sorts of counts. First, in Tables 4 and 5 we provide counts for the five most frequent adjectives that occur in this context with “find” and “consider”, respectively.

	difficult	interesting	easier	necessary	impossible
FIND it _	2325	1028	786	656	506

Table 4. Number of occurrences in COCA of the five most frequent adjectives directly following FIND “it”.

	important	necessary	possible	essential	appropriate, unlikely (tie)
CONSIDER it _	50	42	25	20	19

Table 5. Number of occurrences in COCA of the five most frequent adjectives directly following CONSIDER “it”.

Second, we looked at the frequencies in the same contexts of a sample of adjectives that have been repeatedly classified as PPTs in the philosophical and linguistics literature discussed in section 2, specifically “boring”, “delicious”, “disgusting”, “exciting”, “fun”, and “tasty”. The results appear in Table 6.

	boring	delicious	disgusting	exciting	fun	tasty
FIND it _	67	11	38	49	53	3
CONSIDER it _	0	0	2	0	2	0

Table 6. Number of occurrences in COCA of a sample of predicates of personal taste directly following FIND “it” and CONSIDER “it”.

We now turn to some observations on the results of these searches.

3.3. Observations on the corpus data

First, from the counts in Tables 2 and 3, we can certainly conclude that some moral adjectives, as well as the adverbs “morally” and “ethically”, do appear with “find”. A sample example for each adjective/adverb is provided in (9).

(9) a. [Senator E. Kennedy]: As a matter of your own individual and personal moral beliefs, do you believe that abortion is moral or immoral?
[Judge Souter]: Senator, I'm going respectfully to ask to decline to answer that question for this reason, that whether **I do or do not find it moral** or immoral will play absolutely no role in any decision I make if I am asked to make it on the question of what weight should or legitimately may be given to the interest which is represented by the abortion decision

b. I would never vote for **something I find immoral or unjust** even if 90% of my voters were for it
d. I agree that if I use and enjoy open source software, it is ethical for me to contribute back, and **I find it most ethical to contribute in a fashion that can be used and enjoyed by all those whose contributions I enjoy**
e. [Talking about "pay-to-play" concerts] While **we might find the practice unethical**, disgusting and ugly, it's not illegal
f. In fact, I find it to be a moral responsibility that I take the knowledge that I am able to understand and help make it accessible to everyone. **I find it to be very unethical to provide poor or incomplete information** (which is part of why my posts are so long)
g. T[o] vote for **a third party candidate that I find less morally objectionable** is for me the way to avoid any cooperation with an immoral act
h. I still think even in debate vituperative insults can occasionally be useful, and in less structured discussions elsewhere **I don't find them ethically questionable**, though often overutilized

That said, the data in Tables 2-6 also clearly indicate that our sample moral adjectives and adverbs are as a whole used considerably less frequently with "find" than are the adjectives in our sample of PPTs, both in absolute and relative terms. The PPTs occur vastly more often after FIND "it" than after CONSIDER "it" (where their presence is virtually testimonial). In contrast, the moral adjectives, with the exception of "ethical", appear more often after CONSIDER "it" than after FIND "it" (though overall the numbers are very small), and Table 2 clearly shows a greater tendency to appear as a complement to "consider" than to "find", including for "ethical". Interestingly, however, this asymmetry is not found with the adverbs: indeed, "morally" appears more often as a modifier of an adjectival complement to "find" than it does with adjectival complements to "consider", while "ethically" appears a similar number of times with both.

Of course, these numbers have to be evaluated against the background of other frequency information. The overall frequency of any word will obviously influence how often it appears with other words. Moreover, some words are relatively unselective about the other words they appear with (for example, "be"), while others occur much more frequently with some words than others (such as "radiocarbon", with "dating"). This selectivity, or strength of association, can occur for multiple reasons, both grammatical ("be" is a verb used in a wide range of constructions) and semantic/pragmatic ("radiocarbon dating" describes a particularly widely used method for dating objects, and we may talk considerably less infrequently about other uses of radioactive isotopes of carbon and thus use "radiocarbon" infrequently as a modifier of other terms). In the case that interests us here, it would be interesting to know whether, indeed, there are distinctly different strengths of association between "find" and PPTs, on the one hand, and moral adjectives and "consider", on the other.

In corpus linguistics, one standard measure of strength of association is mutual information (MI), and COCA conveniently provides automatically calculated MI scores word pairs in the

corpus.[4] We will not go into the technical details of MI here (see Evert 2009 for very useful discussion), other than to note that one important limitation of the way in which MI is calculated in COCA is that it does not take into account the syntactic relations between words. It simply looks at cooccurrences within a specified window. Thus, a string like "find the moral responsibility", which is irrelevant for our purposes, contributes to the MI score for "find" and "moral" in exactly the same way as the relevant (if grammatically incomplete) string "find it moral and". In general, the higher a (positive) MI score, the stronger the (positive) association. In the english-corpora.org interface, the default suggestion for a MI search is to find scores of at least 2.5; Hunston (2002, p. 71) asserts that MI scores of "3 or higher can be taken to be significant."

We searched COCA's frequency database for MI scores over 1 for the different adjectives and adverbs mentioned above in combination with "find" and "consider", not placing any minimum threshold on absolute frequencies for the co-occurrences. We found MI scores over 1 for "ethical" (1.96), "unethical" (3.55), and "immoral" (3.63) and "consider"; indeed, "consider" was the verb with the strongest mutual information score for "unethical" and "immoral". Similarly positive scores were found for "morally" (2.29) and "ethically" (2.41) with "consider". "Moral" did not give a positive result in this search. None of the PPTs showed positive MI scores with "consider", and none of the adjectives or adverbs at all showed an MI score over 1 with find, except for "boring" (1.37). Thus, despite the limitations of the MI scores as calculated in COCA, we have some reason to think that moral adjectives and adverbs are semantically different in some way from PPTs, despite the fact that both occur with "find". We did a further search for the adjectives most strongly associated with "find"; the top 10 were "hard-pressed" (4.84), "distasteful" (4.52), "objectionable" (4.38), "amusing" (4.34), "gainful" (3.84), "guilty" (3.79), "off-putting" (3.76), "repulsive" (3.63), "abhorrent" (3.62), and "repugnant" (3.62). Among these, all but "hard-pressed", "gainful", and "guilty" are adjectives that imply an experiential subject, and, to that extent, are arguably PPTs - even if "objectionable", "repulsive", "abhorrent" and "repugnant" can be used for the purpose of assessing moral actions.[5]

---

[4] Mutual Information is calculated in COCA as in (i), taken from https://www.english-corpora.org/mutualInformation.asp with minor modifications.

(i) MI = log((AB*sizeCorpus)/(A*B*span)) / log(2), where
A = frequency of the word of interest (e.g. "moral")
B = frequency of collocate (e.g. "find")
AB = frequency of collocate near the node word (e.g. "find" near "moral")
sizeCorpus= the number of words in the corpus
span = span of words (in COCA, this is 3 to left and 3 to right of word of interest)
log(2) is literally the $\log_{10}$ of the number 2: .30103

[5]Note that in "(not) guilty", "find" often occurs not as a subjective attitude verb but rather acquires a legal sense describing a jury officially deciding on an accused individual's guilt.

4. Discussion

The present work lies within a broader philosophical enterprise of understanding the nature of morality. More precisely, we wish to know to which extent moral judgments are subjective in the way in which judgments of personal taste are. We approach this question by studying moral language; specifically, by looking at how paradigmatic moral predicates – "(im)moral" and "(un)ethical" – combine with subjective attitudes verbs. Overall, our findings show that moral predicates exhibit certain linguistic hallmarks of subjectivity, but, at the same time, behave differently from PPTs.

Both "consider" and "find" (in one of its senses) are verbs that express subjective attitudes. Lasersohn observes that "consider" "is much more limited than 'believe' in the types of complement clause it may combine with. It combines quite naturally with clauses expressing personal taste, but normally does not combine with clauses expressing completely objective matters of fact" (2009, p. 365). "Find", as amply discussed in the literature, clearly tracks subjective judgment and is more restrictive than "consider", since it does not accept complements such as "vegetarian", which are subjective only to the extent that different speakers may appeal to different criteria in classifying things as vegetarian or not.

Our corpus study provides evidence of natural occurrences of moral adjectives with both verbs, so one may be tempted to simply conclude that moral adjectives are therefore subjective, just like PPTs. But this would be a hasty and oversimplified conclusion. Our findings show that moral predicates prefer to occur with "consider" rather than "find" (despite occasionally occurring with the latter), whereas in the case of PPTs, it is the other way around. We draw this more nuanced conclusion from the fact that, as can be seen from Table 2, moral predicates are about three times more likely to occur with "consider" than with "find".[6] While this already reveals a proportional preference for "consider" over "find", it bears noting that the preference is actually much higher, given that the verb "find" is much more frequent than "consider". Table 3, which looks specifically at "find"/"consider it" ADJ constructions, points to a similar pattern. That is to say, while the number of occurrences is altogether low, the much greater frequency of "find" over "consider" suggests that, proportionally, moral predicates show a preference for the latter over the former. On the other hand, as Table 6 shows, PPTs are largely absent in the "consider it" ADJ construction, but very frequent in the "find it" ADJ construction (for instance, the antonyms "fun" and "boring" have over fifty occurrences each, while among moral predicates, "immoral" and "ethical" score highest, with only six occurrences each). Our findings are therefore completely in line with the observation from McNally and Stojanovic (2017), mentioned earlier, that evaluative adjectives (of which moral adjectives are a subtype) occur more naturally with "consider" rather than "find".

[6] Note that this does not extend to predicates that are modified with "morally" (where we see the reverse pattern) and "ethically" (equally likely to occur with either verb). In the corpus data, we often see "find" embed adjectives such as "reprehensible", "objectionable" and "repugnant" modified by "morally". We believe that it is these adjectives that are driving the preference for "find" over "consider", while the adverb "morally" primarily serves to endow the adjectives with a more specific sense.

The corpus study presented here has significant implications for theoretical research on subjectivity. It is generally assumed that "find" is more restrictive than "consider", which in turn is more restrictive than "believe". But the picture that emerges from our corpus search appears to be more subtle. For one thing, pace Lasersohn, "consider" does not "combine quite naturally with clauses expressing personal taste"; witness the fact that predicates such as "boring", "fun", "delicious", "disgusting", and "tasty" hardly ever occur with "consider". This suggests that the relationship between the subjective attitudes that are expressed with the two verbs is not one of subordination. Rather, the attitudes that the two verbs express, considering and finding, may be plausibly seen as involving different types of subjectivity.

A further question is what kind of theory accounts best for the data observed. While there have been a number of interesting and plausible proposals concerning the semantics of "find" (see e.g. Willer, to appear, for references and overview), the question of how it differs from other subjective attitude verbs has been less discussed. A notable exception is the work of Kennedy and Willer (2016, 2022), whose proposal is largely driven by the motivation of capturing the differences between "find" and "consider". They start by outlining their account of attributions of subjective attitudes, based on the idea of what they call "counterstance contingency": "a subjective attitude ascription asserts belief in the proposition expressed by the complement clause, and presupposes the contingency of this belief across a set of contextually provided alternatives to the attitude holder's doxastic state, all of which agree on the salient facts of the matter but disagree on judgments about those facts. We label these alternatives *counterstances* and the contingency across them *counterstance contingency*" (2022, p. 13). After motivating the idea of counterstance contingency, they note: "It remains to explain the more fine-grained differences between *consider*-type and *find*-type subjective attitude verbs. Our key proposal is that the latter presuppose a distinguished kind of subjectivity that we label radical counterstance contingency, which flows from a distinguished kind of pragmatic underdetermination (...)" (*ibid.*, p. 15). While mere counterstance contingency tends to result from incidental underdetermination, radical counterstance contingency results from essential underdetermination. In other words, in the former case, speakers can avoid underdetermination by stipulating that terms be understood in one way rather than another, while in the latter case, their views and experiences diverge more radically and cannot be brought into agreement by mere stipulation.

Kennedy and Willer's account can explain why terms such as "vegetarian" can occur with "consider" but not with "find". It can also provide a plausible story as to why PPTs are more likely to occur with "find" instead of "consider".[7] However, the predictions of their view, when it comes to moral judgments, are less clear. On the one hand, whether a belief is counterstance contingent, and whether it is radically so or not, is conversation-dependent, which would fit well with the observation that moral predicates can occur both in "consider" and in "find" constructions. On the other, this suggests that the interpretation of the attributions of "find"-attitudes vs. "consider"-attitudes should differ precisely along these lines; that is to say, a speaker who uses "consider" presupposes that their (or the

[7] The idea would be, roughly, that if one assumes that judgments of personal taste systematically involve radical counterstance contingency, then speakers should preferably use a verb that triggers this presupposition (to wit, "find") rather than a verb such as "consider", which triggers a weaker presupposition.

attributee's) divergence on moral issues could be settled by stipulation, whereas if they use "find", they presuppose a more radical divergence. It also suggests that the preference for "consider" that we have observed for moral predicates should mirror the tendency of incidental rather than essential underdetermination when it comes to moral judgments. Whether the two predictions are borne out would require examining the examples closely, including a qualitative analysis of the contexts in which they occurred. However, this important task lies beyond the scope of this paper.

5. Conclusion

While the nature of morality has been a core topic of interest for decades, the twenty-first century marks what may be called the empirical turn in philosophy in general, and moral philosophy and philosophy of language in particular. However, the empirical methods used so far have predominantly involved controlled studies involving, for example, the elicitation of acceptability judgments. The present chapter offers new insights on the nature of moral judgments based on corpus methodology. We have presented a corpus study that investigates the subjective character of moral predicates, by examining how they combine with two subjective attitude verbs, "find" and "consider". The study shows, in a nutshell, that moral predicates can occur naturally with both verbs. Nevertheless, they show a clear preference for "consider" over "find". In this respect, moral predicates differ significantly from predicates of personal taste, which embed frequently and naturally under "find", but hardly ever under "consider". This suggests, in turn, that the subjectivity that one sees in moral judgments may well be of a different kind than the subjectivity of personal taste.

Acknowledgments

We would like to thank David Bordonaba, Christopher Kennedy, Malte Willer and an anonymous reviewer for comments. Isidora Stojanovic acknowledges support from the COST Action CA17132 in the Horizon 2020 Framework Programme of the European Union, and from the Agence nationale de la recherche (ANR) under the grant agreement ANR-17-EURE-0017 FrontCog. Louise McNally acknowledges support through an ICREA Academia award.

References

Bordonaba, David. 2017. *Operadores de orden superior y predicados de gusto: Una aproximación expresivista*. Dissertation, Universidad de Granada.

Beebe, James R., and David Sackris. 2016. Moral objectivism across the lifespan. *Philosophical Psychology* 29:912–929.

Björkholm, Stina. 2022. *The duality of moral Language. On hybrid theories in metaethics.* Dissertation, University of Stockholm.

Camp, Elisabeth. 2017. Metaethical expressivism. In *The Routledge handbook of metaethics*, ed. Tristram McPherson and David Plunkett, 87–101. New York and London: Routledge.

Davies, Mark. 2008–. *The Corpus of Contemporary American English (COCA)*. https://www.english-corpora.org/coca/.

Evert, Stefan. 2009. 58. Corpora and collocations. In *Corpus Linguistics: An International Handbook, Volume 2*, ed. Anke Lüdeling and Merja Kytö, 1212–1248. Berlin and New York: De Gruyter Mouton.

Fleischer, Nicholas. 2013. The dynamics of subjectivity. *Proceedings of Semantics and Linguistic Theory* 23:276–294.

Franzén, Nils. 2020. Evaluative discourse and affective states of mind. *Mind* 129:1095–1126.

Goodwin, Geoffrey P., and John M. Darley. 2008. The psychology of meta-ethics: Exploring objectivism. *Cognition* 106:1339–1366.

Goodwin, Geoffrey P., and John M. Darley. 2010. The perceived objectivity of ethical beliefs: Psychological findings and implications for public policy. *Review of Philosophy and Psychology* 1:1–28.

Goodwin, Geoffrey P., and John M. Darley. 2012. Why are some moral beliefs perceived to be more objective than others? *Journal of Experimental Social Psychology* 48:250–256. https://doi.org/10.1016/j.jesp.2011.08.006.

Hunston, Susan. 2002. *Corpora in applied linguistics*. Cambridge: Cambridge University Press.

Karczewska, Natalia. 2019. *Faultless disagreement in contemporary semantic theories*. Dissertation, University of Warsaw.

Kennedy, Christopher. 2013. Two sources of subjectivity: Qualitative assessment and dimensional uncertainty. *Inquiry* 56:258–277. https://doi.org/10.1080/0020174X.2013.784483.

Kennedy, Christopher, and Malte Willer. 2016. Subjective attitudes and counterstance contingency. *Proceedings of Semantics and Linguistic Theory* 26:913–933.

Kennedy, Christopher, and Malte Willer. 2022. Familiarity inferences, subjective attitudes, and counterstance contingency: Toward a pragmatic theory of subjective meaning. *Linguistics and Philosophy.* https://doi.org/10.1007/s10988-022-09358-x

Lasersohn, Peter. 2005. Context dependence, disagreement, and predicates of personal taste. *Linguistics and Philosophy* 28:643–686.

Lasersohn, Peter. 2009. Relative truth, speaker commitment and control of implicit arguments. *Synthese* 166:359–374. https://doi.org/10.1007/s11229-007-9280-8.

McNally, Louise, and Isidora Stojanovic. 2017. Aesthetic adjectives. In *The semantics of aesthetic judgment*, ed. James Young, 17–37. Oxford: Oxford University Press.

Odrowąż-Sypniewska, Joanna. 2021. Vagueness in natural language. In *The Cambridge handbook of the philosophy of language*, ed. Piotr Stalmaszczyk, 434–449. Cambridge: Cambridge University Press.

Pölzler, Thomas. 2017. Revisiting folk moral realism. *Review of Philosophy and Psychology* 8:455–476.

Pölzler, Thomas, and Jennifer Cole Wright. 2020a. An empirical argument against moral non-cognitivism. *Inquiry*. https://doi.org/10.1080/0020174X.2020.1798280.

Pölzler, Thomas, and Jennifer Cole Wright. 2020b. Anti-realist pluralism: A new approach to folk metaethics. *Review of Philosophy and Psychology* 11:53–82.

Railton, Peter. 2017. Naturalistic realism in metaethics. In *The Routledge handbook of metaethics*, ed. Tristram McPherson and David Plunkett, 43–57. New York and London: Routledge.

Reuter, Kevin, Lucien Baumgarten, and Pascale Willemsen. (ms.). Tracing thick and thin concepts through corpora. Unpublished manuscript.

Rudolph, Rachel. 2021. Talking about appearances: The roles of evaluation and experience in disagreement. *Philosophical Studies* 177:197–217.

Sarkissian, Hagop. 2016. Aspects of folk morality: Objectivism and relativism. In *A companion to experimental philosophy*, ed. Wesley Buckwalter and Justin Systma, 212– 224. Malden and Oxford: Wiley Blackwell. https://doi.org/10.1002/9781118661666.ch14.

Sayre-McCord, Geoff. 2005. Moral realism. In *The Stanford encyclopedia of philosophy*, ed. Edward N. Zalta, Summer 2021 edition. https://plato.stanford.edu/archives/sum2021/entries/moral-realism/.

Sæbø, Kjell Johan. 2009. Judgment ascriptions. *Linguistics and Philosophy* 32:327–352.

Silk, Alex. 2021. Evaluational adjectives. *Philosophy and Phenomenological Research* 102:127–161.

Solt, Stephanie. 2018. Multidimensionality, subjectivity and scales: Experimental evidence. In *The semantics of gradability, vagueness, and scale structure*, ed. Elena Castroviejo, Louise McNally, and Galit Weidman Sassoon, 59–91. Cham: Springer. https://doi.org/10.1007/978-3-319-77791-7_3.

Soria Ruiz, Andrés, Bianca Cepollaro, and Isidora Stojanovic. 2021. The semantics and pragmatics of value judgments. In *The Cambridge handbook of the philosophy of language*, ed. Piotr Stalmaszczyk, 434–449. Cambridge: Cambridge University Press.

Soria Ruiz, Andrés, and Federico Faroldi. 2020. Moral adjectives, judge-dependency and holistic multidimensionality. *Inquiry* 64:1–30.

Stephenson, Tamina. 2007. Judge-dependence, epistemic modals, and predicates of personal taste. *Linguistics and Philosophy* 30:487–525.

Stojanovic, Isidora. 2007. Talking about taste: Disagreement, implicit arguments, and relative truth. *Linguistics and Philosophy* 30:691–706.

Stojanovic, Isidora. 2012. Emotional disagreement. *Dialogue* 51/1: 99–117.

Stojanovic, Isidora. 2017. Context and disagreement. *Cadernos de Estudos Lingüísticos* 59:9–22.

Stojanovic, Isidora. 2019. Disagreements about taste vs. disagreements about moral issues. *American Philosophical Quarterly* 56: 29–42.

Stojanovic, Isidora. to appear. Evaluativity. In *The Oxford handbook of contemporary philosophy of language*, ed. Ernest Lepore and Una Stojnic. Oxford: Oxford University Press.

Stojanovic, Isidora, and Elsi Kaiser. 2022. Exploring valence in judgments of taste. In *Perspectives on taste*, ed. Jeremy Wyatt, Julia Zakkou, and Dan Zeman, 231–259. New York: Routledge.

Sundell, Timothy. 2011. Disagreements about taste. *Philosophical Studies* 155:267–288.

Umbach, Carla. 2021. Evaluative predicates. Beyond *fun* and *tasty*. In *The Wiley Blackwell companion to semantics*, ed. Daniel Gutzmann, Lisa Matthewson, Cecile Meier, Hotze Rullmann, and Thomas Ede Zimmermann. https://doi.org/10.1002/9781118788516.sem127.

Väyrynen, Pekka. 2013. *The rude, the lewd and the nasty: A study of thick concepts in ethics*. Oxford: Oxford University Press.

Verheyen, Steven, Sabrina Dewil, and Paul Egré. 2018. Subjectivity in gradable adjectives: The case of *tall* and *heavy. Mind and Language.* https://doi.org/10.1111/mila.12184.

Willer, Malte. to appear. Subjectivity. In Ernest Lepore and Una Stojnic (Eds.), *The Oxford handbook of contemporary philosophy of language*. Oxford: Oxford University Press.

Woods, Jack. 2017. The Frege-Geach Problem. In *The Routledge handbook of metaethics*, ed. Tristram McPherson and David Plunkett, 226–242. New York and London: Routledge.

Wright, Jennifer Cole, Piper Grandjean, and Cullen McWhite. 2013. The meta-ethical grounding of our moral beliefs: Evidence for meta-ethical pluralism. *Philosophical Psychology*, 26: 336-361

Zakkou, Julia. 2019. *Faultless disagreement: A defense of contextualism in the realm of personal taste*. Frankfurt am Main: Klostermann.

Zeman, Dan. 2020. Faultless disagreement. In *The Routledge handbook of philosophy of relativism*, ed. Martin Kusch, 486–495. New York and London: Routledge.