



ARTICLE

On Being Ignorant

Anne Meylan¹  and Kevin Reuter² 

¹Department of Philosophy, University of Zurich, Zürich, Switzerland and ACEPS, Johannesburg, University of Johannesburg and ²Department of Philosophy, Linguistics and Theory of Science, University of Gothenburg, Gothenburg, Sweden
Corresponding author: Anne Meylan; Email: anne.meylan@uzh.ch

Abstract

A traditional view holds that ignorance is simply the absence of knowledge. This view has recently been challenged by the Normative Account, which sees ignorance as involving a normative failure. In this paper, we argue that both perspectives capture important insights. Drawing on three empirical studies, we propose a two-senses account of ignorance, according to which there exist two notions of ignorance: one normative, the other non-normative. We also offer a new explanation of the normative aspect. Our findings suggest that what is negative in being ignorant lies in the expectation that one ought to know the relevant fact.

Keywords: ignorance; knowledge; normativity; empirically-informed conceptual analysis; empirical studies; corpus analysis

1. Introduction

What is ignorance? An influential and intuitive view is the following:

The Knowledge Account: A subject *S* is ignorant of the fact that *p* iff *S* does not know that *p*.¹²

The Knowledge Account is reflected in some dictionaries' definitions. According to the Oxford English Dictionary, for instance, to be ignorant is to be

destitute of knowledge, either in general or with respect to a particular fact or subject.

This article provides empirical evidence for the view that the Knowledge Account cannot fully explain what ignorance is. More precisely, our empirical studies (Studies 1 and 2) show that *ignorance*³ (or, at least one notion of it⁴) has a normative dimension that is left unexplained

¹The main contender of the Knowledge Account is the True Belief Account: A subject *S* is ignorant of the fact that *p* iff *S* does not hold the true belief that *p* (Peels, 2010, 2011, 2012). We say more about the True Belief Account, and whether the results of our empirical studies concern this account too, in the final discussion below.

²As the definition above makes clear, we assume in this article that we are ignorant only of true propositions. We cannot be ignorant of falsehoods. This is a widely shared view. Nevertheless, Le Morvan (2022) has defended the idea that ignorance can also concern falsehoods. In particular, he shows that it is possible to be ignorant of *p* when *p* is false. For example, it seems that we can be ignorant of propositions that are necessarily false (e.g., that 193,846,592,857 is a prime number) because we have never considered them. For reasons of space, we limit ourselves in this article to examining the normativity of the more "classical" ignorance that concerns true propositions.

³In this article, we follow the convention to italicize when we speak about the concept of *ignorance*, and use quotation marks to indicate the word "ignorance."

⁴This qualifier is important. A detailed explanation of why we have to temper our conclusion in this way will be clear once we discuss the results of our empirical studies.

by the Knowledge Account. We then propose a novel explanation of the nature of the normative dimension of *ignorance*. With the help of Study 3, we substantiate the hypothesis that what is normatively negative in ignorance has to do with our expectations regarding what one should know: what is negative in being ignorant about p is that one is expected to know that p .

What the results of our studies additionally show is that the concept of *ignorance* is not “all of a piece.” Rather, there seems to be both a normative and a non-normative notion of *ignorance*. We conclude from this that the most suitable account of ignorance is a two-senses account.

2. Theoretical Background

As just mentioned, the Knowledge Account is a very influential view of ignorance. Recently, however, it has come under attack. Pritchard (2021a, 2021b) has offered several counterexamples that suggest *ignorance* has a normative dimension, which the Knowledge Account cannot easily capture (see also Meylan 2020, 2024).⁵ Let us point out that we present these counterexamples because one of our objectives in this article is to consolidate the (possibly refutable) intuitions underlying them, using empirical data. So, here are Pritchard’s three counterexamples:

First, in Pritchard’s view, it is quite unfitting to attribute ignorance of a fact to individuals when this fact cannot possibly be known. For instance, it does not sound fully appropriate to claim that “prehistorians are ignorant of whether *Homo sapiens* sapiens were tying their hair up.” We would rather say that they simply do not know this, or that they simply have no belief about this.

Second, according to Pritchard, it seems quite incongruous to attribute ignorance of a fact to individuals when knowing this fact is pointless, and this even when it is possible to get to know it (with some tenacity, perhaps). For instance, it does not sound fully appropriate to tell someone (while on holiday in Normandy): “You are ignorant of the number of grains of sand there are on the D-Day beach.” It seems more appropriate to say: “You do not know how many grains of sand there are on the D-Day beach.”

A third kind of counterexample to the Knowledge Account and the True Belief Account consists of cases in which the attribution of ignorance of a fact to an individual seems inappropriate because of the strong moral reasons that speak against this individual getting to know this fact. Here is Pritchard’s persuasive example:

Consider, for example, a scrupulous District Attorney who realizes, when looking through the files she has been given as part of a case, that she has been handed information that was meant for the defence team rather than for the prosecution. She might be fully aware that this file contains information that would have an important bearing on the case; indeed, which might even decisively indicate the defendant’s guilt. Moreover, she might know that the facts in this file would not ordinarily be otherwise epistemically available to her, such that this is effectively her one chance to know them. Nonetheless, being scrupulous, she declines to open the file and returns it to the defence team unopened. There are thus truths that were directly epistemically available to the District Attorney but which she lacks knowledge of as a result of her conscious decision to lack them, based on good moral reasons. Significantly, however, if one were in full possession of why she lacks this knowledge, then I don’t think there would be any temptation to describe this lack of knowledge as ignorance. (2021b, 8)

⁵The reason why the Knowledge Account cannot easily capture this normative aspect is that the mere lack of knowledge does not necessarily entail a normative failure. Even though knowledge might be a valuable mental state, the absence of this valuable mental state does not necessarily consist in a failure, in something negative. The same is analogously true for the valuable state of pleasure: the absence of pleasure does not necessarily consist in displeasure (see Chisholm & Sosa, 1966). It might well consist in having neither displeasure nor pleasure, that is, in something neither good nor bad.

So, to reiterate, three types of counterexamples contradict the Knowledge Account according to Pritchard.⁶ The absence of knowledge does not intuitively qualify as an instance of ignorance:

- When an individual does not know a fact that she cannot know.
- When an individual does not know a fact that would be pointless to know.
- When an individual does not know a fact that would be morally problematic to know.

This leads Pritchard to suggest an alternative account of ignorance, according to which ignorance does not merely consist in the absence of knowledge but also encloses a normative dimension.⁷ Following Pritchard, we shall call this alternative account the Normative Account of ignorance.

Normative Account

An individual *S* is ignorant of the fact that *p* iff:

- (i) *S* is lacking a certain cognitive state (e.g., she does not know, does not hold the true belief that *p*, etc.) (= the non-normative condition); and
- (ii) *S* ought (for one reason or another) to be in this cognitive state (= the normative condition).

As it should be clear from the foregoing discussion, Pritchard's objection relies heavily on intuitions regarding specific cases and will not impress those who do not share his intuitions. In view of this, the main purpose of this article is to provide some additional credibility to the Normative Account by running three empirical studies that suggest *ignorance* is (often)⁸ understood to have a negative normative dimension. Before presenting these studies, three remarks are in order.

First, there are several theoretical advantages to acknowledging the normative dimension of *ignorance* (as we are about to do below). Mainly, if ignorance is the lack of knowledge that one ought to have, one can explain why the individuals involved in the three kinds of counterexamples above are not ignorant, even though they do not know. Indeed, going in the reverse order, it is, for moral reasons, not the case that the District Attorney ought to obtain knowledge about the content of the file. Because it is pointless to know how many grains of sand there are on D-Day beach, it is not the case that the Normandy holiday-goer ought to know this. And it is not the case that the prehistorians ought to know whether *H. sapiens* sapiens were tying their hair up. This results from the "ought implies can principle": since the prehistorians cannot know this, this is not something they ought to know.

As Pritchard emphasizes (2021b, 11–12), another theoretical advantage of the Normative Account is that it re-establishes a normative resemblance between:

- (a) the character trait that one attributes with the predicate "ignorant" (as when one says "he is an ignorant man") and;
- (b) the ignorance that is mainly at issue in this article, namely, the cognitive state in which an individual is when she is ignorant of a fact.

⁶In his 2021a paper, Pritchard presents a fourth kind of counterexample to the Knowledge Account. Wittgensteinian hinge propositions—such as the proposition that I have two hands or that the earth is not going to blow up in 2 s—are propositions that we do not know since, per definition, they fall outside of the range of propositions for which we can provide reasons and, thereby, out of the range of propositions that we can know. However, we are not ignorant of hinge propositions. We are, in fact, certain of their truth.

⁷See also Meylan, 2020, 2024 for an alternative defense of the normative conception of ignorance.

⁸Here again, the qualifier is important. This is because, as we discuss in detail later, we do not deny that there is a non-normative notion of ignorance too.

When ignorance is understood as a character trait, ignorance seems to enclose a negative normative dimension.⁹ Given the plausibly close conceptual relation between character-ignorance and state-ignorance, it seems rather strange that they would differ so much with respect to their normative dimension, that only character-ignorance is characterized by a negative normative dimension. The Normative Account alleviates this worry by attributing a negative normative dimension to state-ignorance as well, thereby re-establishing a normative resemblance between the two.

The second remark is a clarification. Even though we will end up defending the view that *ignorance*—or at least one key notion of ignorance—has a negative normative dimension, we do not deny that ignorance can be bliss in certain contexts. It is important in this regard to distinguish the *prima facie* necessary normative dimension of *ignorance* from the all-things-considered contingent normative features of specific instances of ignorance. For instance, suppose that, while you are struggling to submit an article on time, you get to know in advance as a result of an indiscretion that your grant application has been rejected for the second time. Since this news has devastating consequences for your work motivation, it would certainly have been better if you had remained ignorant of your failure. In this case, ignorance is all-things-considered better than knowledge in virtue of certain contingent features that this instance of ignorance possesses, namely, it would avoid the loss of your work motivation. But saying this is perfectly compatible with the Normative Account. What the Normative Account states, to be sure, is that ignorance has a *prima facie* negative normative dimension and that this *prima facie* negative dimension is constitutive of ignorance. The Normative Account does not deny that this *prima facie* negative normative dimension can be overridden by contingent normative features in many cases.

As for the third remark, we would like to recall that the Knowledge Account is not the only non-normative understanding of ignorance on the market. The Knowledge Account has a main contender, the True Belief Account, according to which:

A subject *S* is ignorant of the fact that *p* iff *S* does not hold the true belief that *p* (Peels, 2010, 2011, 2012, 2023).

Due to space constraints, we cannot present additional studies that would target the True Belief Account more specifically. But there are sound reasons to think that such studies would replicate the results that we present below. We explain why this is so in the final discussion of our results (see Section 5.c).

3. The Normative Dimension of *Ignorance*

In the previous section, we saw that Pritchard primarily relies on intuitions from thought experiments to challenge the Knowledge Account of ignorance. While we do not believe his thought experiments to be particularly problematic, empirical studies have shown that such intuitions often fail to be robust and representative of the respective ordinary concepts that philosophers try to capture (e.g., Beisbart & Reuter, 2021; Wiegmann, Horvath, & Meyer, 2020). We, therefore, decided to run three empirical studies. In this section, we present the results of a corpus study as well as a vignette study, designed to trace the normative dimension of “ignorant” and, thereby, to investigate whether the Normative Account is supported by empirical data.

3.a A corpus analysis on “ignorant” and “unaware”

Following Firth’s dictum that “You shall know a word by the company it keeps!” (Firth, 1957, 11), we think one of the most promising strategies to reveal a normative component of “ignorant,” is to

⁹When understood as a character trait, it seems quite natural to identify ignorance as an intellectual vice, that is, as a normatively negative disposition that is susceptible to manifest itself in several reprehensible cognitive behaviours, such as the lack of proper care and attention to certain facts.

analyze the *company* that the word “ignorant” keeps. If we were to find evidence of a predominantly normative company, this would indicate a predominantly normative use, and, in turn, would suggest that “ignorance” has a normative dimension that the Knowledge Account cannot easily capture.

To reveal a possible normative dimension, we decided to operationalize the extent to which “ignorant” is used normatively by investigating those adjectives that are connected with the adjective “ignorant” using the conjunction “and.” The conjunction “and” is often used to conjoin adjectives that have the same polarity (Elhadad & McKeown, 1990). This means that if “ignorant” has a negative normative dimension, we would predict that people conjoin other normative terms of the same polarity with our target adjective. Thus, “arrogant and ignorant” and “lazy and ignorant” should be typical and very common, whereas “easygoing and ignorant” and “content and ignorant” should be less common and atypical. This approach has also been recently employed by Willemssen, Baumgartner, Frohofer, and Reuter (2023) and Reuter, Baumgartner, and Messerli (2025).

Having this operationalization at hand, how do we now go about collecting data that speaks in favor of or against the normative use of “ignorant”? Philosophers have increasingly turned to large corpora to investigate the usage of terms and phrases, and make use of the sophisticated tools that linguists and computer scientists have developed to analyze large amounts of data that are compiled into corpora.¹⁰ Some of these corpora are freely available online, e.g., the News on the Web (NOW) corpus, one of the largest open-access corpora, containing over 20 billion words from web-based newspapers and magazines from 2010 to the present (as of January 2025).

When we insert “ADJ and ignorant of” into the search field of the NOW corpus, we find that the most frequent adjectives conjoined with “and ignorant of” are “naïve,” “insensitive,” “stupid,” and “illiterate.” While this result provides some initial evidence of a normative use of “ignorant,” the data are far too sparse. Thus, we decided to build our own corpus composed of comments from the social media website Reddit.

Methods and hypothesis. Importantly, we only collected data in which the term “ignorant” is likely to be used in an epistemic sense, but not as an attribution of character. To realize this restriction, all the contexts featuring “ignorant” are those in which “ignorant” is followed by the preposition “of.”¹¹ Thus, examples included “This is a fallacy and frankly *arrogant and ignorant of* the psychological research surrounding it” and “They were *lazy and ignorant of* environmental pollution” (our italics). We also needed to include a control condition with which to compare the use of “ignorant of.” Such a control condition is necessary to rule out possible confounding factors that might have distorted the results of our study. Ideally, we would have liked to collect data that feature the phrase “[Adj] and does not know of *p*,” given that the knowledge account proposes that these statements are synonymous or near-synonymous. However, “does not know” allows for constructions that are different from “ignorant.”¹² This might lead to all sorts of unwanted effects. We, therefore, decided to collect texts in which people used the phrase “[Adj] and unaware of *p*.” Indeed, as has been pointed out (Silva & Siscoe, 2023, 11) ignorant and un/aware have “similar semantic constructions whether or not their objects are facts.” Furthermore, substituting “does not

¹⁰Within the last 15 years, philosophers have begun utilizing corpus analytic tools to investigate philosophical questions about pain (Reuter, 2011), appearance language (Fischer, Engelhardt, & Herbelot, 2015), Nietzsche’s use of “drive” and “instinct” (Alfano, 2018), knowledge (Hansen, Porter, & Francis, 2021), causation (Sytsma, Bluhm, Willemssen, & Reuter, 2019), conspiracy theory (Reuter & Baumgartner 2025), and many more. For an overview, see Reuter & Baumgartner, 2024, as well as the forthcoming edited volume on *Advances in Experimental Philosophy and Corpus Methods* by Ulatowski, Weijers, & Sytsma *forthcoming*.

¹¹An alternative strategy would have been to search for the phrase “ignorant that.” However, many instances of “ignorant that” do not specify the object of ignorance but are instead followed by a result clause, such as “my father was so ignorant that I could hardly stand to be around him.” To avoid these character-attribution uses, we searched only for the phrase “ignorant of.”

¹²For instance, Le Morvan (2015) has argued that the complement of ignorance of a proposition *p* is not knowledge that *p* but knowledge of *p*.




Figure 1. Sentiment values of various adjectives from the dictionary *SentiWords*.

know p ” by “unaware of p ” in the test condition is harmless for another reason. There are several ways of describing what it means for an individual to be aware of a fact. Two possibilities mentioned by Silva and Siscoe are as follows:

- An individual is aware of a fact p if he is in a position to know that p ; and
- An individual is aware of a fact p if that individual represents the fact in question, and the representation he makes of p is “suitably non-accidental.”

These are two rival conceptions, but we do not have to take sides. We need only note that, according to both theories, awareness is a necessary condition of knowledge (for more details about this, see Silva & Siscoe, 2023, 15). There is, therefore, an absence of knowledge in every case of unawareness. This is why the substitution mentioned above is not problematic. Even if we use “unaware” rather than “does not know,” unawareness implies the absence of knowledge.

As we are primarily interested in the way laypeople use the term “ignorant,” a corpus using posts and discussion notes from Reddit, a social news website, was built (no specific subreddits were selected).¹³ We then searched for and collected comments featuring the phrases “[Adj] and ignorant of p ,” and “[Adj] and unaware of p .” Comments that include a negation of the adjectives (e.g., “not,” “hardly,” or “barely”) or any other adverbial modifier (e.g., “very,” “rather,” or “mostly”) were discarded. Our final corpus consisted of 1,626 comments for “[Adj] and ignorant of p ” and 1,748 comments for “[Adj] and unaware of p .”

To measure the evaluative intensity of the conjoined adjectives, we used the sentiment dictionary *sentiWords*. *sentiWords* contains sentiment values, i.e., values that indicate both the polarity and the strength of the polarity for tens of thousands of English words. Values range from “−1,” meaning highly negative, to “+1,” meaning highly positive. The evaluative intensity of the target term would then be determined by calculating the average sentiment value from all conjoined adjectives. To give the reader a better feel for the range of values that are assigned to various words, Figure 1 displays some exemplary adjectives (figure taken from Willemsen et al., 2023).

We are now in a position to state the primary hypothesis of the corpus study.

Hypothesis 1 The average evaluative intensity of “ignorant” is significantly more negative than the evaluative intensity of the control term “unaware.”

Results. Figure 2 shows the sentiment distribution as well as the means and median values for both target terms. The evaluative intensity for “ignorant” was −0.26; for the control term “unaware,” the evaluative intensity was calculated to be −0.14. As the sentiment values of the conjoined adjectives are not normally distributed, we used a one-sided unpaired two-sample Wilcoxon test ($W = 2,482,818$, p -value < 0.001), which shows that “ignorant” has significantly lower conjoined sentiment values, on average. Consequently, Hypothesis 1 cannot be rejected.

We also decided to have a more qualitative look at the data. Table 1 below lists the 20 most frequent adjectives that are conjoined with “ignorant of p ” as well as “unaware of p .” These

¹³We would like to thank Lucien Baumgartner for his support in building the corpus and running the analysis on the data.

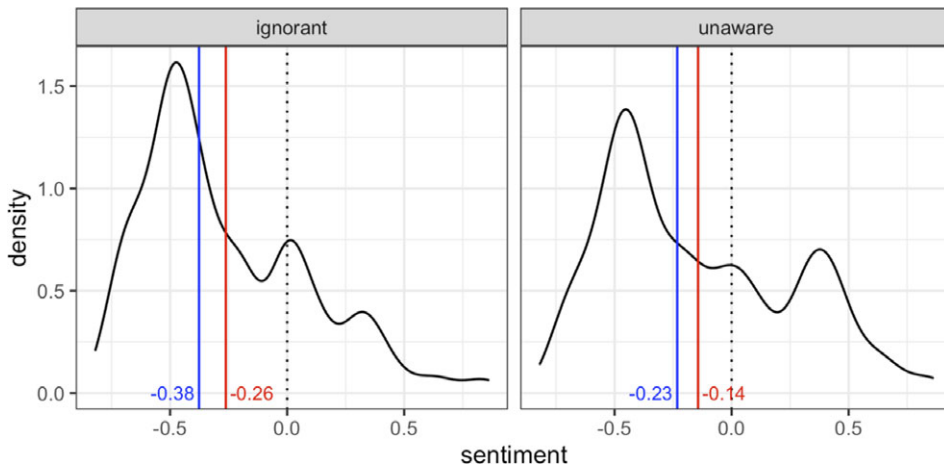


Figure 2. Sentiment distribution. The panels show the distribution of the sentiment values for the terms “ignorant” and “unaware.” The red line depicts the mean, the blue one the median.

adjectives cover 45.1% (resp. 44.5%) of all uses we collected and, thus, while not comprehensive, provide a representative picture of the data.

Discussion. Our data analysis reveals a strong co-occurrence of negatively connoted adjectives with the phrase “and ignorant of *p*.” The only two neutral terms among the top 20 terms are “young” and “innocent.” In contrast, almost half (9 out of 20) of the 20 most frequently conjoined terms with “and unaware of *p*” are descriptive or neutral terms. Furthermore, while negative terms that co-occur with “unaware” are almost exclusively epistemic (the only exception is “selfish”), many terms (6 out of 20) that occur together with “ignorant” are morally thick terms like “selfish,” “arrogant,” “lazy,” “rude,” and “disrespectful.” Defenders of the knowledge account might highlight that both target terms are often conjoined with negatively connoted epistemic terms like “stupid,” “dumb,” and “naïve.” Note, however, that “stupid” and “naïve” are conjoined roughly twice as often with “ignorant of *p*” compared to “unaware of *p*.”

While almost all of the 20 most frequent adjectives conjoined with “and ignorant of” are primarily negative, and the mean sentiment value is both significantly below the midpoint as well as significantly below the mean value for “unaware,” one might object that a sentiment value of -0.26 is not *strongly* negative. Wouldn’t we expect an even more negative value? To put our data into perspective, Table 2 displays the results for a range of adjectives for which Willemsen et al. (2023) collected data using a similar method.¹⁴ A comparison with the other terms suggests that “ignorant” is grouped with other clearly negative epistemic terms such as “inconsistent,” and “illogical.”

3.b Cancellability task

The results from the corpus study above suggest the term “ignorant” to be a normative term. However, the data are admittedly only indicative, as we have not directly tapped into the putative normative dimension of the term. Corpus studies, for the most part, only provide a record of a term’s use but do not manipulate any variable directly. Thus, given that it would be desirable to have more direct evidence for the normative dimension of “ignorant,” we decided to run an empirical study using the cancellability test (see, e.g., Grice, 1975; Zakkou, 2018).

The cancellability test is often employed to examine whether information is only conversationally implicated or likely to be semantically entailed. Take the following statement: “This is a lake, but

¹⁴The major difference is that Willemsen et al. (2023) searched for phrases of the form “[Adj] and [target term],” whereas in our case we needed to include the preposition “of” and, thus, searched for “[Adj] and ignorant of.”

Table 1. A list of the 20 most frequent adjectives conjoined with “and ignorant of” and “and unaware of”

x and ignorant of			x and unaware of		
Term	Number	Percentage	Term	Number	Percentage
Stupid	116	7.13	Ignorant	103	5.89
Naive	87	5.35	Young	89	5.09
Dumb	59	3.63	Stupid	59	3.38
Selfish	53	3.26	Dumb	58	3.32
Young	43	2.64	Naive	46	2.63
Short-sighted	42	2.58	Asleep	44	2.52
Wrong	42	2.58	Unskilled	44	2.52
Arrogant	41	2.52	Aware	40	2.29
(Close-, narrow-, or simple-) minded	36	2.21	New	38	2.17
Uneducated	29	1.78	Innocent	29	1.66
Racist	26	1.60	Confused	27	1.54
Blind	20	1.23	Uneducated	27	1.54
Lazy	20	1.23	Selfish	26	1.49
Dismissive	18	1.11	Unconscious	24	1.37
Rude	18	1.11	Blind	23	1.32
Unaware	18	1.11	Stationary	22	1.26
Disrespectful	17	1.05	Incompetent	21	1.20
Innocent	17	1.05	Happy	20	1.14
Uninformed	17	1.05	Old	19	1.09
Delusional	14	0.86	Delusional	18	1.03
	733	45.08		777	44.45

We color-coded the terms according to the following criteria (the coding was done by the authors; as such, there is bound to be some disagreement on which adjectives belong in which category): yellow indicates a negative epistemic term; green indicates a negative non-epistemic term; and blue indicates the counterpart term (unaware vs. ignorant). If the term is not color-coded, then the term is primarily descriptive.

by that I am not saying that it consists of water.” This statement is contradictory and is indeed considered by people to be contradictory. Now, consider the statement: “This chocolate is good value-for-money, but by that I am not saying that we should buy it.” While a speaker might indeed express her wish to buy that chocolate when uttering “This chocolate is good value-for-money,” such information is only conversationally implicated, and not semantically entailed. Consequently, it is not contradictory to deny that one should buy the chocolate.

Willemssen and Reuter (2021) have started to use the cancellability test to examine the evaluative dimension of thick terms. For instance, they presented participants with statements such as “John is rude, but I don’t mean to say something negative about him.” Most people rate this statement to be strongly contradictory, which suggests that the negative evaluation is not merely pragmatically conveyed but semantically entailed. Given these previous results, we decided to adapt the

Table 2. Sentiment values for a list of epistemic and legal thick terms, as well as descriptive terms

Category	Term	Value
Epistemic	Inappropriate	−0.40
Epistemic	Illogical	−0.33
Epistemic	Inconsistent	−0.27
Target term	Ignorant	−0.26
Descriptive	Complex	−0.01
Descriptive	Active	0.07
Descriptive	Limited	0.09

experimental paradigm to investigate whether knowledge and ignorance attributions behave differently with respect to the cancellability of a normative component.

Methods. We randomly assigned 172 participants ($M_{age} = 36.33$ years ($SD = 11.40$), 85 females, 86 males, 1 other), who we recruited on *Prolific*, to one of the following four statements:

- (Bees Ignorance) “Tom is ignorant of the fact that bees produce honey, but by that I am not saying that he should be aware of this fact.”
- (H₂O Ignorance) “Tom is ignorant of the fact that water consists of H₂O, but by that I am not saying that he should be aware of this fact.”
- (Bees Knowledge) “Tom does not know that bees produce honey, but by that I am not saying that he should be aware of this fact.”
- (H₂O Knowledge) “Tom does not know that water consists of H₂O, but by that I am not saying that he should be aware of this fact.”

The two (ignorance) statements served as our test condition, while the two (knowledge) statements were part of the control condition. We chose the phrase “ignorant of the fact that” as it is the most frequent (according to the NOW corpus) and a very natural way to utilize “ignorant” in the epistemic sense. The latter part of all four statements, i.e., “but by that I am not saying that he should be aware of this fact,” was always the same.¹⁵ By stating that latter part, the speaker intends to cancel the normative content “should be aware of *p*” that might be part of the meaning of “is ignorant of *p*.” After the participants had read the statement, they were asked to rate how much they thought that the speaker, Sally, contradicted herself on a 9-point Likert scale anchored at “1” = “definitely not” till “9” = “definitely yes.” High ratings suggest that the participant believes the normative component “should be aware of *p*” to be semantically entailed by the first part of the statement. Low ratings indicate that the participant considers the normative content to be at most conversationally implicated. Our main hypothesis for this study was as follows:

Hypothesis 2 Contradiction ratings are higher for ignorance statements compared to knowledge statements.¹⁶

Results. The mean rating for the knowledge statements was 2.13 ($SD = 2.00$), indicating that participants thought that statements of the form “S does not know that *p*” do not express a

¹⁵There are, of course, other normative contents that might be constitutive of the normative dimension of “ignorance” that we did not test for.

¹⁶The study’s hypothesis and criteria were preregistered at the Open Science Framework. The dataset of this study is available in [this Open Repository](#).

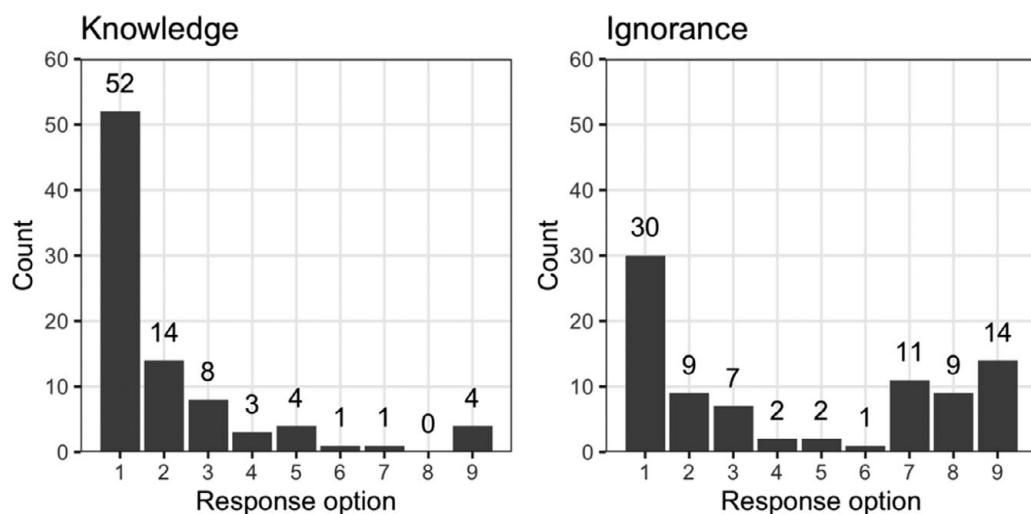


Figure 3. Response profiles for the knowledge (left) and the ignorance condition (right). The scale was anchored at 1 = “definitely not” and 9 = “definitely yes”.

normative component. In contrast, the mean value for ignorant statements was 4.33 (SD = 3.28), which was significantly ($p < 0.001$) different from the control condition, suggesting that the phrase “is ignorant of p ” is considered to be (at least more strongly) normative. We conducted a two-way ANOVA with *Epistemic Term* (ignorance vs. knowledge) and *Scenario* (bees vs. H₂O) as independent factors, and *Contradiction Ratings* as the dependent variable. *Epistemic Term* was a significant factor ($F(1, 168) = 28.43, p < 0.001, \eta^2 = 0.15$). This effect reveals that participants rated the statements as more contradictory when they included the phrase “is ignorant of p ” compared to “does not know that p .”

We were also interested in finding out whether the results show either a greater inclination of participants to rate the ignorance statements as more contradictory or a split in their ratings. In the latter case, we should find that some participants consider the ignorance statements highly contradictory, whereas others would rate the statements as not contradictory. Figure 3 shows the distribution of the responses for the data from the knowledge and ignorance statements. The graph on the right reveals that the responses for the ignorance condition are strongly bimodally distributed.

Discussion. The results of Study 2 provide further evidence for a normative reading of “ignorant.” The average contradiction ratings in Study 2 were not only significantly higher in the test condition compared to the control condition but also fairly close to the midpoint of “5.” To examine whether participants find ignorance statements only “more strongly” contradictory, we checked the distribution of the responses. Perhaps surprisingly, participants seemed to be split between a descriptive and a normative interpretation of the statement in question.¹⁷ There are various ways to account for bimodally distributed data. We will discuss these options further in the General Discussion (Section 5) when we present our own account.

4. Expectations and Normativity

Study 1 and Study 2 provide empirical evidence that *ignorance* is taken (at least by many) as a normative notion. As such, these studies put some pressure on the Knowledge Account. As already

¹⁷According to the standard definition of *contradiction*, a statement is either a contradiction or it is not. The bimodal distribution might calm worries about using a graded scale when asking a contradiction question.

discussed in Section 2, this is because the Knowledge Account cannot easily account for the normative dimension of *ignorance*.

Study 3 below does two things. It puts additional pressure on the Knowledge Account by showing that there are circumstances in which people do not find it fitting to attribute ignorance to an individual even though the latter fails to know a piece of information. But Study 3 also casts light on the very nature of the normative dimension that *ignorance* seems to entail. With Study 3, we aim to investigate which normative factor drives people's attribution of ignorance. Glimpsing at the wealth of data from our corpus analysis, a very common pattern emerged. People often claimed that another person was ignorant of p in cases in which they *expected* that person to know p .¹⁸ We rarely found statements in which such ignorance was attributed in circumstances in which we would generally not expect that person to know. Given these qualitative observations, we state the following hypothesis:

Hypothesis 3 The more expected a proposition p is perceived to be known, the more fitting it is to claim that subject S is ignorant of the fact that p , in case S does not know that p .

4.a Methods

To test Hypothesis 3, we presented participants with either a statement that people likely expect another person to know (high expectation content) or a statement that people would not expect another person to know (low expectation content).¹⁹ All subjects were randomly assigned to a single vignette that had the following structure.

S is a middle-aged American citizen who has lived in the United States for all her/his life.

p . S is not aware of the fact that p .

(e.g., Bees produce honey. Tom is not aware of the fact that bees produce honey.)

How fitting is the following claim:

S is ignorant of the fact [$/$ does not know] that p .

Half of the participants read the ignorance claim (" S is ignorant of the fact that p ."), whereas the other half were presented with the knowledge claim (" S does not know that p ."). All participants were then asked to rate the fittingness of the claim on a scale from "1" = not at all fitting to "7" = absolutely fitting. The presented propositional contents p are all true, were selected from two different categories (biology and chemistry), and were related to those used in Study 2 for congruency.

High expectation content:

- Bees produce honey.
- Water consists of H_2O molecules.

¹⁸The notion of expectation used here expresses a form of minimal normativity that is not necessarily associated with stronger normative concepts such as obligations, blame, and so forth. We employ this notion in its ordinary sense, and in this ordinary sense, expectation is tied to a weak form of requirement. For instance, when we say that X is expected to attend the party, we mean that X is supposed to attend. This does not necessarily imply that X has an obligation to attend or would deserve blame if they do not. Similarly, when we say that X is expected to know that p , we mean that X is supposed to know that p (based, for example, on their background knowledge).

¹⁹Study 3 was also preregistered at the Open Science Framework, and the dataset of this study is available in this Open Repository.




Figure 4. Average fittingness ratings in the four tested conditions. Error bars indicate standard error around the means.

Low expectation content:

- Some species of ants produce honey.
- Fructose consists of $C_6H_{12}O_6$ molecules.

4.b Results

A total of 277 participants ($M_{age} = 38.39$ years ($SD = 12.76$), 138 females, 136 males, and 3 others) were recruited on *Prolific*. The mean fittingness ratings are displayed in Figure 4. We performed a two-way ANOVA with the independent factors *Expectation* (high vs. low) and *Epistemic term* (ignorance/knowledge), as well as the dependent variable *Fittingness*. Whereas the independent factor *Epistemic term* was highly significant, $F(273) = 41.486$, $p < 0.001$, $\eta^2 = 0.132$, the factor *Expectation* was not, $F(273) = 1.067$, $p = 0.303$. Importantly, the interaction between *Epistemic term* and *Expectation* was highly significant, $F(273) = 19.324$, $p < 0.001$, $\eta^2 = 0.066$. Accordingly, Hypothesis 3 cannot be rejected.

4.c Discussion

In the knowledge condition, the participants considered the knowledge claim (*S* does not know that *p*.) to be highly fitting, as was expected. The ratings were also not significantly dependent on whether the content was expected to be known. In contrast, in the ignorance condition, there was a significant difference between the high-expectation sentences and the low-expectation sentences. In those cases, in which the content was expected to be known, the subjects thought the ignorance claim to be highly fitting. However, when the content was not expected to be known, the mean rating was significantly reduced.

The results of Study 3 reveal that the less we expect a person to be aware of a fact, the less likely people are willing to ascribe ignorance of the fact that *p* to that person who fails to know *p*, providing strong support for our hypothesis. In other words, the outcome suggests that the attribution of ignorance is strongly driven by the expectations we have of people to know a certain fact. The results of Study 3 also reveal that there are cases in which a person does not know *p* but at the same time is not considered to be ignorant of *p*. In other words, in some circumstances, not knowing that *p* is not a sufficient condition for being ignorant of *p*.

A similar question arises for the results of Study 3 that we already discussed in Study 2. Do people merely feel less inclined to ascribe ignorance of a fact to a person when the fact is less expected to be




Figure 5. Response profile for the ignorance condition (low expectation content). The scale was anchored at 1 = “not at all fitting” and 7 = “absolutely fitting”.

known, or are people split between those who consider such an ignorance attribution highly fitting and those who consider the claim to be not fitting at all? Figure 5 below clearly shows a bimodal distribution, similar to the one we found in Study 2. Whereas 41% of the participants gave ratings below the midpoint, 56% of the participants considered the statement to be fitting.

5. General Discussion

In the preceding two sections, we presented the findings from three empirical studies. The upshot of our studies is as follows. First, the results of Study 1 and Study 2 suggest that at least one notion of ignorance (more on this below) has a normative dimension. The results of Study 3 allow us to say more about the normative dimension that characterizes *ignorance*. Simply put, attributions of ignorance are more fitting the more we expect a person to know something. In the concluding section of this article, we interpret our study findings and propose a novel account of ignorance, referred to as the “two senses account” (Section 5.a). In Section 5.b, we consider an objection that could be made to us by the defender of the knowledge account (in short, the objection consists of advocating a pragmatist interpretation of the empirical data obtained). Finally, we briefly discuss the potential extension of our results to the true belief account of ignorance in Section 5.c.

5.a Developing a two-senses account

Our studies unveil the presence of a normative dimension within the concept of ignorance. However, we also find some evidence that people entertain a descriptive reading of “ignorant” that is at least consistent with the knowledge account. Recall that in Study 2, more than half of the participants did not consider it contradictory to claim “Tom is ignorant of p , but by that I am not saying that he should be aware of p .” How should we interpret this result? The response patterns we recorded seem to show that we are dealing with two different but related senses of “ignorant,” one descriptive sense along the lines of the knowledge account, and one normative sense along the lines favored by Pritchard’s view. Given that almost all participants in Study 2 either entertained a descriptive reading or a normative reading, we call this account the two-senses account.

According to the two-senses account, there are two notions of *ignorance*, one normative and one non-normative.

According to the non-normative understanding, *ignorance* is not a normative notion, and it might well consist in the mere lack of knowledge that *p* (since the mere lack of knowledge is not necessarily negative; After all, having no consideration whatsoever about *p* is a way of lacking knowledge that *p* that is neither positive nor negative).

According to the normative understanding, *ignorance* of *p* is the lack of knowledge that *p*, where knowledge that *p* is expected.²⁰

Importantly, further evidence for the two-senses account is provided by Study 3. Our main aim was to investigate the proposal that expectations drive the normative reading of “being ignorant of a fact.” While we were able to present evidence in favor of this proposal, we also reproduced the bimodal response distribution that we found in Study 2. This bimodal distribution bolsters the two-senses account in that it suggests that people either interpret claims including the term “ignorant” in a normative manner or in a purely descriptive manner.

The Knowledge Account cannot fully explain ascriptions of ignorance. Why? Because our empirical studies show that there are two notions of *ignorance*, one normative, one non-normative, and the Knowledge Account is not such a two-senses account. It might well capture the non-normative notion, but it fails to capture the normative one. However, in the next subsection, we will discuss a plausible objection that advocates of the knowledge account might raise (this is the pragmatist objection to which we have alluded before).

Similarly, the normative account supported by Pritchard (2021b) also cannot explain the bimodal distribution of the results of Study 2 and Study 3. Even though people often believe that normative content is semantically entailed by “being ignorant,” this is not always the case. This is something that our two-senses account of *ignorance* fully acknowledges.

In this article, we not only cast doubt on both the Knowledge Account and the Normative Account but also make a suggestion as to what the additional normative dimension that characterizes the normative notion of *ignorance* might be. We suggest that ignorance that *p* is the lack of knowledge that *p*, where knowledge that *p* is expected. This conclusion leads us to three important considerations.

First, it is important to highlight that our approach to understanding the normativity of ignorance provides an explanation as to why the individuals involved in Pritchard’s counterexamples are not considered ignorant. Indeed, none of them are expected to possess knowledge about the things they are unaware of (for instance, the prehistorians are not expected to know the specific hairstyles of *Homo sapiens sapiens*).

Second, the two-senses account possesses an advantage by aligning with the existence of two distinct understandings of *ignorance* within contemporary philosophical discourse. As mentioned previously, one classical understanding of ignorance is non-normative, i.e., ignorance as the lack of knowledge or the lack of true belief. But another, much more normatively loaded notion of ignorance is becoming increasingly influential in the philosophical literature. Most strikingly, perhaps, Charles Mills’s concept of *white ignorance* (2007) is normative. In our view, there is no opposition between these two views. The fact that the philosophical literature contains these two notions is simply mirroring the truth of the two-senses account we end up defending on the basis of our experimental studies.

Third, let us emphasize that we do not consider our explanation of the normativity of *ignorance* to be the only plausible one. We believe philosophical work on the concept of *ignorance* has only just begun, as many questions remain unanswered. There are certainly other factors—besides expectations of knowledge—that can influence attributions of ignorance and deserve to be discovered.

²⁰Several authors, such as Notelmann (2016) and Peels (2023), distinguish between various kinds of ignorance but do not go as far as stating that the notion has two senses.

Other intriguing questions are, “Do our results only hold for the English language or also for a wide variety of languages and cultures?” and “Is “ignorant” a thick term?”

Before concluding, let us address a counterargument that proponents of the knowledge account might put forth. Subsequently, we will explore how our findings might extend to the true belief account of ignorance.

5.b A pragmatist solution for the knowledge account and why it fails

The Knowledge Account holds that a person is ignorant of p if and only if that person does not know that p . Prima facie, our data suggest that the Knowledge Account cannot be the whole story. While not knowing p might indeed be a necessary condition for being ignorant of p , it does not seem to be sufficient. At least for many, a normative condition must be satisfied for ascriptions of “ignorant” to be appropriate.

Is there any way for a proponent of the Knowledge Account to explain the data from Studies 1, 2, and 3? We believe the most plausible strategy for advocates of the Knowledge Account to preserve a descriptive reading is to explain the normative content as being pragmatically communicated and not as being part of the semantics of “ignorant.” Careful readers of the empirical sections of this article might think, “But wasn’t Study 2 supposed to rule out a pragmatic reading of the normative component?” Indeed, the cancellability test is designed to find out whether certain content is at most conversationally implicated. However, the pragmatist might respond in saying that particularized conversationally implicated content, e.g., the chocolate case, is not the only pragmatic content in town. Generalized conversationally implicated content might be harder to cancel.

Conversational implicatures fall roughly into two groups. First, particularized conversational implicatures are activated in specific circumstances. In contrast, generalized conversational implicatures are less dependent on the context in which they are uttered. Consider the chocolate case from above: “This chocolate is good value-for-money.” The conversational implicature “Let us buy this chocolate” might be triggered when being in a store looking at that chocolate. However, if you have already bought the chocolate and have just enjoyed parts of it, the implicature is likely not to be communicated.

Here are two generalized conversational implicatures that are less dependent on context:

- Generalized 1: “Jenny tried to get into the club, but by that I am not saying that she failed to get into the club.”
- Generalized 2: “I washed my hands and I had dinner, but by that I am not saying that I washed my hands before I had dinner.”

While the term “try” implicates failure in a wide variety of contexts, phrases of the term “[action1] and [action2]” almost always implicate that [action1] was or will be performed first and [action2] afterwards. Fortunately, we know the response distributions from previous studies (Willemsen & Reuter, 2021) on such generalized implicatures. The response profile for these two conditions was markedly different from the response profile we observed in Study 2 (Section 3.b), with almost everyone considering these cases to be highly contradictory. Thus, if the normative content of “being ignorant of p ” is communicated pragmatically, it can hardly be a case of a generalized conversational implicature.

There is a second reason why we are skeptical that a proponent of the Knowledge Account can easily go pragmatist in regards to the normative content. If the normative content were communicated pragmatically, but was not part of the semantics of the term, then fittingness ratings in Study 3 should display a unimodal distribution. At a minimum, we would expect that everybody *feels* the normative pull of “ignorant” but to different degrees, some stronger, and some weaker. But that is not what we found. Instead, we found a bimodal distribution, which, when adopting a pragmatist solution, would lead to the rather strange result that only some people understand the pragmatic implicature, while others do not.

In summary, we do not believe that a pragmatist solution can be defended. In other words, we do not believe that the Knowledge Account can be argued to provide a unified account of the meaning of “ignorant.”

5.c *The true belief account in light of the results*

As mentioned above, the main contender of the Knowledge Account is the True Belief Account defended by Peels (2010, 2011, 2012, 2023) in which ignorance consists in the absence of true belief. A legitimate question at this point is whether the studies presented above let us predict anything regarding the truth or falsity of the True Belief Account.

The results of Study 1 (the Corpus Analysis) might well be extended to the True Belief Account. Recall that our main hypothesis is that the average evaluative intensity of “ignorant” is significantly more negative than the evaluative intensity of the control term “unaware.” Since lacking true belief about a fact is a way of being unaware of it, the results of the Corpus Analysis are also valid for the True Belief Account.

We can wonder whether the cancellability task will give us similar results if the control conditions were:

- (Bees True Belief) “Tom does not hold the true belief that bees produce honey, but by that I am not saying that he should be aware of this fact.” and
- (H₂O True Belief) “Tom does not hold the true belief that water consists of H₂O, but by that I am not saying that he should be aware of this fact.”

It would be interesting to run a study including these control conditions and to compare the results to the ones we obtained with our own Study 2. Even though it is difficult to predict what the outcome of such a study would be, a plausible hypothesis is that the contradiction rating might end up being even lower for (Bees True Belief) and (H₂O True Belief) than with (Bees Knowledge) and (H₂O Knowledge). After all, we often attribute false beliefs to people about a fact without meaning that they should be aware of that fact. It does not seem that the attribution of a lack of true belief semantically entails the fact that this belief should have been held. But, of course, this hypothesis would need to be verified empirically.²¹

We might also wonder whether Study 3 would give us similar results if we had instead asked people whether they find it fitting to attribute “S does not have a true belief that *p*” in (i) low or (ii) high-expectation scenarios. A reasonable hypothesis is that the level of expectation would not have any substantial influence over how people evaluate the fittingness of attributing a lack of true belief (just as it does not have any substantial influence over how people evaluate the fittingness of attributing a lack of knowledge). It is probable that people would find it very fitting to attribute a lack of true belief that *p* to an individual when *S* lacks awareness of the fact that *p*, and this, independently, of whether *S* was expected to hold the true belief that *p* or not. To think otherwise would require us to doubt that people understand the lack of awareness of the fact that *p* as involving the absence of the belief that *p*, which does not seem very reasonable to doubt.

6. Conclusion

Our studies uncover that ignorance encompasses both a normative and a non-normative dimension. The traditional knowledge account falls short in capturing the normative aspect observed in our studies, while Pritchard’s normative view fails to account for the descriptive dimension. In

²¹Indeed, if this hypothesis were to be empirically confirmed, it would not confer any advantage upon the True Belief Account, as it would demonstrate that this account likewise encounters significant difficulties in accounting for the variation in response distribution in the test conditions.

response, we propose a comprehensive framework—the two-senses account—that adeptly addresses people’s perspectives on ignorance. Moreover, we highlight the substantial role of expectations in attributions of ignorance. Nevertheless, further research is needed to examine additional factors that shape attributions of ignorance in diverse contexts, languages, and cultures.

Acknowledgments. We are very grateful to Lucien Baumgartner for his invaluable assistance in constructing the Reddit corpus and conducting the corpus-data analysis. We also thank two anonymous reviewers of this journal for their helpful and constructive feedback. This research was presented at the Workshop on Ignorance in Regensburg, the Experimental Philosophy Conference in Granada, the GAP conference in Berlin, and research colloquia in Bern, Lund, and Zurich, and we are grateful to the attendees for their thoughtful questions and comments. Special thanks go to Manuel Gustavo Isaac, Ethan Landes, Thomas Raleigh, and Pascale Willemsen for their feedback on an earlier draft of this article.

Funding statement. This work was supported by the Swiss National Science Foundation under grant numbers 100012_189259 (A.M.) and PCEFP1_181082 (K.R.).

Anne Meylan is a Professor of Theoretical Philosophy at the University of Zurich. Her research primarily concerns the rationality of belief and the normative structures that govern intellectual life, with a particular focus on the phenomena of suspension of judgment and ignorance.

Kevin Reuter is a senior lecturer at the University of Gothenburg. His research interests are in the Philosophy of Mind, Language, Decision Making, Truth and Conspiracy Theories, as well as Experimental Philosophy.

References

- Alfano, M. (2018). Digital humanities for history of philosophy: A case study on Nietzsche. In I. Levenberg, L. Neilson, T. Rheams, D. (eds) *Research Methods for the Digital Humanities* (pp. 85–101). Cham: Palgrave Macmillan.
- Beisbart, C., & Reuter, K. (2021). What is the folk concept of life?. *Australasian Journal of Philosophy*, 101(2), 486–507.
- Chisholm, R., & Sosa, E. (1966). On the logic of intrinsically better. *American Philosophical Quarterly*, 3(3), 244–249.
- Elhadad, M., & McKeown, K. (1990). Generating connectives. *Proceedings of the 13th conference on Computational linguistics*, 3, 97–101.
- Firth, J. R. (1957). A Synopsis of Linguistic Theory 1930–1955. *Studies in Linguistic Analysis*, 1–32.
- Fischer, E., Engelhardt, P., & Herbelot, A. (2015). Intuitions and illusions: From explanation and experiment to assessment. In E. Fischer & J. Collins (Eds.), *Experimental philosophy, rationalism, and naturalism* (pp. 267–300). London and New York: Routledge.
- Grice, H. (1975). Logic and conversation. In P. Cole & J. L. Morgan (Eds.) *Syntax and Semantics*, Vol. 3, Speech Acts, (pp. 41–58), New York: Academic Press.
- Hansen, N., Porter, J. D., & Francis, K. (2021). A corpus study of “know”: On the verification of philosophers’ frequency claims about language. *Episteme*, 18(2), 242–268.
- Le Morvan, P. (2015). On the ignorance, knowledge, and nature of propositions. *Synthese*, 192, 3647–3662.
- Le Morvan, P. (2022). Ignorance, truth, and falsehood. *Ratio*, 35(3), 169–180.
- Meylan, A. (2020). Ignorance and its disvalue. *Grazer Philosophische Studien*, 93(3), 433–447.
- Meylan, A. (2024). In defence of the normative account of ignorance. *Erkenntnis*, 89(1), 207–221.
- Notelmann, N. (2016). The varieties of ignorance. In R. Peel and M. Blaauw (Eds.), *The epistemic dimensions of ignorance*. Cambridge: Cambridge University Press. (pp. 33–56).
- Peels, R. (2010). What is ignorance? *Philosophia*, 38(1), 57–67.
- Peels, R. (2011). Ignorance is lack of true belief. *Philosophia*, 39(2), 345–355.
- Peels, R. (2012). The new view on ignorance undefeated. *Philosophia*, 40(4), 741–750.
- Peels, R. (2023). *Ignorance: A philosophical study*. Oxford: Oxford University Press
- Pritchard, D. (2021a). Ignorance and inquiry. *American Philosophical Quarterly*, 58(2), 111–124.
- Pritchard, D. (2021b). Ignorance and normativity. *Philosophical Topics*, 49(2), 225–244.
- Reuter, K. (2011). Distinguishing the appearance from the reality of pain. *Journal of Consciousness Studies*, 18(9–10), 94–109.
- Reuter, K., & Baumgartner, L. (2024). Corpus analysis — Building and using corpora: A case study on the use of ‘conspiracy theory’. In S. Kornmesser, M. Bauer, et al. (Eds.), *Experimental philosophers for beginner*. Switzerland: Springer.
- Reuter, K., & Baumgartner, L. (2025). Conspiracy theories are not theories: Time to rename conspiracy theories. In M. Isaac, S. Koch, & K. Scharp (Eds.), *New perspectives on conceptual engineering*. Synthese Library: Springer.
- Reuter, K., Baumgartner, L., & Messerli, M. (2025). A corpus study on the normativity of rationality. *Thought*.
- Silva, P., & Siscoe, R. W. (2023). Ignorance and awareness. *Noûs*, 58(1), 225–243.

- Sytsma, J., Bluhm, R., Willemsen, P., & Reuter, K. (2019). Causal attributions and corpus analysis. In E. Fischer, & M. Curtis (Eds.), *Methodological advances in experimental philosophy*, London: Bloomsbury Academic. (pp. 209–238).
- Ulatowski, J., Weijers, D., & Sytsma, J. (forthcoming). *Advances in experimental philosophy and corpus methods*. London: Bloomsbury Academic.
- Wiegmann, A., Horvath, J., & Meyer, K. (2020). Intuitive expertise and irrelevant options. *Oxford Studies in Experimental Philosophy*, 3(3), 275.
- Willemsen, P., & Reuter, K. (2021). Separating the evaluative from the descriptive: An empirical study of thick concepts, *Thought: A Journal of Philosophy*, 10(2), 135–146. <https://doi.org/10.1002/tht3.488>
- Willemsen, P., Baumgartner, L., Frohofer, S., & Reuter, K. (2023). Examining evaluativity in legal discourse: A comparative corpus-linguistic study of thick concepts. In *Advances in experimental philosophy of law*, London: Bloomsbury Academic.
- Zakkou, J. (2018). The cancellability test for conversational implicatures. *Philosophy Compass*, 13(12), e12552.