

What isn't obvious about 'obvious': A data-driven approach to philosophy of logic

Moti Mizrahi

Florida Institute of Technology

Forthcoming in A. Aberdein and M. Inglis (eds.), *Advances in Experimental Philosophy of Logic and Mathematics* (London: Bloomsbury Press).

Abstract: It is often said that 'every logical truth is obvious' (Quine 1970: 82), that the 'axioms and rules of logic are true in an obvious way' (Murawski 2014: 87), or that 'logic is a theory of the obvious' (Sher 1999: 207). In this chapter, I set out to test empirically how the idea that logic is obvious is reflected in the scholarly work of logicians and philosophers of logic. My approach is data-driven. That is to say, I propose that systematically searching for patterns of usage in databases of scholarly works, such as JSTOR, can provide new insights into the ways in which the idea that logic is obvious is reflected in logical and philosophical practice, i.e., in the arguments that logicians and philosophers of logic actually make in their published work.

Keywords: argument; indicator words; logic; methodology; obvious; philosophical logic; philosophy of logic

1. Introduction

What kinds of insights can we gain by analyzing a corpus of logical and philosophical texts? In this chapter, I explore one way in which the methods of data science can be used to analyze a

large corpus of logical and philosophical texts in order to shed light on logical, philosophical, and metaphilosophical questions. I illustrate this data-driven approach with an example from logic and the philosophy of logic. In particular, I set out to test the way in which the oft-repeated thesis that ‘logic is obvious’ is reflected in logical and philosophical practice. I hypothesize that, if logicians and philosophers of logic subscribe to the view that logic is obvious, then we should expect to find them using ‘obvious’ in contexts in which they signal the certainty of the claims and/or arguments they make. If they use ‘obvious’ when their claims and/or arguments betray a lack of certainty, then that would suggest that their use of ‘obvious’ should not be taken literally and that their view, as evinced by what they do (more precisely, the arguments they put forward), is that logic is not obvious.¹

To put it another way, I think that the text mining, analysis, and visualization techniques of data science can help us address the following question: What would logic, philosophical logic, and philosophy of logic look like *in practice* if practitioners subscribed to the view that logic is obvious? Instead of picking out a few examples of logicians and philosophers of logic who have asserted that logic itself, or some proper subset of logic, is obvious, and deriving a general lesson from such examples, I examine a large corpus of text and look for correlations that are designed to pick out such assertions on a much larger scale. My hope is that this data-driven methodology can be used to investigate systematically other claims that are often made about fields like logic and philosophy as well as other disciplines.²

¹ Attempts to incorporate empirical or experimental methods into philosophy have often been met with the “that’s not philosophy” (Jenkins 2014) or “that’s not philosophically significant” charge. (See, e.g., Kauppinen 2007. Cf. Knobe 2007; O’Neill and Machery 2014.) This issue is beyond the scope of this chapter. For present purposes, I hope it is enough to point out that the volume for which this chapter was invited is about empirical or experimental approaches to logic and/or philosophy of logic and mathematics. In that respect, this chapter offers an empirical approach to a question about logical and philosophical practice.

² See, e.g., Ashton and Mizrahi (2018) for an empirical, data-driven investigation of how the idea that philosophy is a priori is reflected in philosophical practice.

For the purposes of this empirical study, I have focused on the idea that logic is obvious. Many logicians and philosophers of logic have expressed approval of the thesis that logic is obvious. In a recent post on the blog of the American Philosophical Association, Steven M. Chan (2017) relates the following story about Willard van Orman Quine:

Once in [an introduction to symbolic logic] course, after he wrote a proof on the board, a student raised his hand and asked impatiently, “Why bother writing out that proof? It’s obvious.” To which Quine replied, “Young man, this entire course is obvious.”

Quine expresses this idea explicitly in print when he says that ‘every logical truth is obvious’ (Quine 1970: 82) and that ‘elementary logic is obvious or can be resolved into obvious steps’ (Quine 1976: 112). Other logicians and philosophers of logic share this view that ‘logic is a theory of the obvious’ (Sher 1999: 207). According to Murawski (2014: 87), the ‘axioms and rules of logic are true in an obvious way,’ and Shenefelt and White (2013: 301) say that ‘a good deal of logic is, in fact, so intuitively obvious that it needs no further justification.’ In addition to being obvious, logic is also said to be ‘necessary and a priori’ (Russell 2015: 793).³ Even those who do not subscribe to the idea that logic is obvious, necessary, and a priori claim that, both in principle and in practice, ‘good students in an intro logic class may regard classical logic, not just as a reasonable theory, but as *obviously correct*’ (Russell 2015: 797; emphasis in original). The point, then, is that the word ‘obvious’ gets used quite frequently when logicians and philosophers of logic describe logic or engage in the philosophy of logic. Moreover, standard logic texts make liberal use of ‘obvious’ in describing proofs, theorems, consequences, axioms, rules, etc. For instance, ‘Double negation is fairly obvious and needs little explanation’ (Hurley 2012: 402).

³ According to Field (1998: 4), ‘it makes sense to regard logic as a priori but at the same time to think it conceivable that further conceptual developments could show it not to be, by showing that logic is empirically defeasible (in interesting ways) after all.’ Cf. Bueno (2010: 105) who argues that fallibilism and apriorism are not reconcilable.

As Sher (2010: 158) points out, however, ‘the idea that logic is obvious is a vague idea.’

The word ‘obvious’ itself is vague in more ways than one. According to the *Oxford English*

Dictionary, ‘obvious’ is used as either an adjective or a noun in the following ways:

1. Plain and evident to the mind; perfectly clear or manifest; plainly distinguishable; clearly visible.
2. Lacking in subtlety, sophistication, or originality; banal, predictable.
3. Natural, likely; such as common sense might suggest.

In addition to these different ways in which ‘obvious’ can be used, various things can be said to be obvious. For example:

1. Propositions can be said to be obvious in the sense of being *evident to the mind*.
Explanations can be said to be obvious in the sense of being *perfectly clear*. Shades of color can be said to be obvious in the sense of being *plainly distinguishable*. Macroscopic objects, like trees and cars, can be said to be obvious in the sense of being *clearly visible* to a sensory apparatus.
2. Punchlines can be said to be obvious in the sense of *lacking in subtlety or sophistication*.
Theses can be said to be obvious in the sense of *lacking in originality*. Consequences can be said to be obvious in the sense of being *predictable*.
3. Events can be said to be obvious in the sense of being *natural* or *likely*. Decisions can be said to be obvious in the sense of being *what common sense might suggest*.

Accordingly, the meaning of ‘obvious’ will depend on the ways in which the word ‘obvious’ is used and the things that are said to be obvious. In other words, *context* partly determines the meaning of ‘obvious’. In this chapter, then, I set out to examine the contexts in which the word ‘obvious’ is used in logic, philosophical logic, and the philosophy of logic. Through a systematic

examination of these contexts, I aim to test empirically how the idea that ‘logic is obvious’ is reflected in logical and philosophical practice. My approach is data-driven. That is to say, I propose that systematically searching for patterns of usage in databases of scholarly works, such as JSTOR, can provide new insights into the ways in which the idea that logic is obvious is reflected in logical and philosophical practice, i.e., in the arguments that logicians and philosophers of logic actually make in their scholarly publications.⁴

2. Methods

Most logic textbooks instruct students to look for indicator words when trying to identify arguments in texts. For example, according to Salmon (2013: 39), indicator words are

Words commonly used to signal premisses or conclusions of arguments. Examples of premiss indicator words are *for*, *since*, *because*, and *for the reason that*. Examples of conclusion indicator words are *hence*, *thus*, *therefore*, and *so*, *it follows that*, and *for that reason* (emphasis in original).⁵

Indicator words are also supposed to help in distinguishing between deductive and inductive arguments. For example, according to Baronett (2016: 23):

To help identify arguments as either *deductive* or *inductive*, one thing we can do is look for key words or phrases. For example, the words “necessarily,” “certainly,” “definitely,” and “absolutely” suggest a deductive argument. [...] On the other hand, the words “probably,” “likely,” “unlikely,” “improbable,” “plausible,” and “implausible” suggest inductive arguments (emphasis added).

⁴ As an anonymous reviewer pointed out, what one considers obvious might sometimes be left unsaid. As we will see, ‘obvious’ talk is quite widespread in logic and philosophy journal articles and is becoming significantly more widespread over the years.

⁵ See also Bessie and Glennan (2000: 4), Copi et al. (2011: 11-12), and Marcus (2018: 10).

Likewise, according to Hurley (2012: 34), ‘inductive indicators’ include terms and phrases such as ‘probably’, ‘improbable’, ‘plausible’, ‘implausible’, ‘likely’, ‘unlikely’, and ‘reasonable to conclude’, whereas ‘deductive indicators’ include terms and phrases such as ‘it necessarily follows that’, ‘certainly’, ‘absolutely’, and ‘definitely’.

Accordingly, I propose to use indicator words to test how the idea that ‘logic is a theory of the obvious’ (Sher 1999: 207) and that ‘logic is [...] necessary and a priori’ (Russell 2015: 793) is reflected in the scholarly work of logicians and philosophers of logic. If logicians and philosophers of logic take logic to be obvious, necessary, and a priori, then the arguments they actually make in practice, i.e., in scholarly publications, should reflect that in some way. More explicitly, my empirical study is designed to address the following questions:

- Does ‘obvious’ correlate with deductive indicators, such as ‘necessarily’ and ‘certainly’?
- Does ‘obvious’ correlate with inductive indicators, such as ‘probably’ and ‘likely’?
- Does ‘obvious’ correlate with hedging markers, such as ‘suggest’ and ‘seem’?

I propose that, if logicians and philosophers of logic accept the idea that logic is obvious, necessary, and a priori, then this idea would be reflected in actual practice, i.e., in the way in which logicians and philosophers of logic *do* and *talk* about logic in argumentative writings, as follows:

(A) If practitioners understand logic as obvious, necessary, and a priori, then we would expect to find positive correlations between the word ‘obvious’ and deductive indicators like ‘necessarily’ and ‘certainly’.

(B) If practitioners understand logic as obvious, necessary, and a priori, then we would expect to find negative correlations (or no correlations at all) between the word ‘obvious’ and inductive indicators like ‘probably’ and ‘likely’.

(C) If practitioners understand logic as obvious, necessary, and a priori, then we would expect to find negative correlations (or no correlations at all) between the word ‘obvious’ and hedging markers like ‘suggest’ and ‘seem’.

In other words, if logicians and philosophers of logic subscribe to the view that logic is obvious, necessary, and a priori, then we would expect to see them make the sort of arguments that reflect that when they do logic or talk about logic. That is to say, we would see instances of ‘obvious’ occur mostly in the context of deductive arguments, rather than inductive arguments, since deductive arguments are said to be ‘indefeasible’ and are the sort of arguments whose premises ‘provide conclusive support for their conclusions’ or ‘necessitate the truth of the conclusion’ (Sinnott-Armstrong and Fogelin 2015: 181).⁶

To be clear, I am not suggesting that logicians and philosophers of logic make only deductive arguments. They probably make all kinds of arguments in their scholarly work. What I am suggesting, however, is that, if one believes that logic is obvious, necessary, and a priori, then that belief should be reflected in the sort of arguments one makes when one does logic or talks about logic. Now, since deductive arguments are the sort of arguments whose premises ‘provide conclusive support for their conclusions’ or ‘necessitate the truth of the conclusion’ (Sinnott-Armstrong and Fogelin 2015: 181), we would expect practitioners who think that logic is obvious, necessary, and a priori, to make such arguments more often than other kinds of arguments (e.g., inductive arguments) when they do logic or talk about logic. In other words, if practitioners subscribe to the view that logic is obvious, necessary, and a priori, then we would

⁶ According to Augustus De Morgan (1839: 3), ‘the question of logic is, does the conclusion certainly follow if the premises be true?’

expect to see more deductive arguments than other kinds of argument in logical and philosophical practice in the context of talk about the obviousness of logic.⁷

The data driving this empirical study is taken from JSTOR Data for Research (jstor.org/dfr). This database allows researchers to search full texts for exact phrases and access the metadata associated with search results. I used this database to search for the term ‘obvious’, as well as ‘obviously’ and ‘obviousness’ (JSTOR allows for truncation or ‘wildcard’ search using ‘obvious*’), through research articles written in English. JSTOR does not have a discipline category for logic in particular, so I have created a dataset from the five logic journals in the JSTOR database, namely, *The Journal of Symbolic Logic* (1936-2012), *Studia Logica* (1953-2012), *Journal of Philosophical Logic* (1972-2012), *Journal of Logic, Language, and Information* (1992-2012), and *The Bulletin of Symbolic Logic* (1995-2012).

Having selected the journals for data mining, I then searched for ‘obvious*’ in the context of deductive indicators, inductive indicators, and hedging markers through research articles contained in my Logic dataset, and ran statistical analyses on search results relative to their proportions in the JSTOR corpus in order to test empirically how the idea that logic is obvious is manifested in logical and philosophical practice. To make sure that the aforementioned indicator words occur in the contexts of arguments, to rule out non-argumentative (e.g., rhetorical) instances of ‘obvious’ as much as possible, and to test predictions (A)-(C) above, I have anchored them to the argument indicators ‘therefore’ and ‘follows’ within 10 words of each other (using the operator ~10 in JSTOR’s search box). This method yields the deductive indicator pairs and inductive indicator pairs listed in Table 1.

⁷ See, for example, Suppes’ (1999 :129) reason for not including mathematical content in the practice exercises for chapter 4 of his logic textbook: ‘it is often intuitively *obvious that the conclusions logically follow from the premises given in the exercises*’ (emphasis added).

Here is an example of ‘obvious’ and the argument indicator ‘therefore’ that this search methodology would pick out:

The possibility of situations with the classical valuation of conjunction, disjunction and conditional is *obvious*. *Therefore ...* (Kovač 2012: 347).

Likewise, here is an example of ‘obvious’ and the argument indicator ‘follows’ that this search methodology would pick out:

Obviously, every admissible numbering x is computable. Moreover, we have that $\beta \leq x$.
Since Z is dense in Q_c , *it follows that* Q_c is recursively separable (Spren 1998: 195).

I have also anchored the hedging markers ‘seem’ and ‘suggest’ to the argument indicators ‘therefore’ and ‘follows’ to make sure that my search methodology will pick out usage of these hedging markers in argumentative contexts.⁸ This method yields the following hedging indicator pairs: ‘therefore seem’, ‘follows seem’, ‘therefore suggest’, and ‘follows suggest’ (see Table 1).

Table 1. Deductive, inductive, and hedging indicator pairs

Deductive indicator pairs	Inductive indicator pairs	Hedging indicator pairs
therefore necessarily	therefore probably	therefore seem
therefore certainly	therefore likely	therefore suggest
follows necessarily	follows probably	follows seem
follows certainly	follows likely	follows suggest

This search methodology will pick out instances of these indicator words such as the following (emphasis added):

⁸ On the hedging uses of ‘seem’ in philosophy, see Cappelen (2012: 46). On both ‘seem’ and ‘suggest’ as hedging markers in corpus linguistics and natural language processing, see Thabet (2018: 679-681).

Obviously, Γ and Π are \mathbf{Z} -specific. [...] either the antecedent is not \mathbf{Z} -specific, or the succedent is not necessarily irreducible (Petersen 2000: 396).

Unlike modal logic [...] and thus the elimination of variables is not *likely* to proceed. [...]

Models can be built over the set C in the *obvious* way (Kracht 2013: 1338).

Partial logic (in which some statements receive no truth value) *suggests* itself as an *obvious* candidate... (Horsten 2015: 687).

For the purposes of comparison, I ran the same searches I did on the Logic dataset on three discipline categories in the JSTOR database: Philosophy, Mathematics, and Biological Sciences. I have selected the Philosophy discipline category on JSTOR because it is useful to compare it to the Logic dataset. Most of the journals contained in JSTOR's Philosophy category publish work in so-called mainstream areas of philosophy, by professionally trained philosophers, that purport to offer arguments, which should, in theory, contain many indicator words. The so-called 'general' philosophy journals, such as *Mind* and *Noûs*, often publish work on logic, philosophical logic, and the philosophy of logic. In fact, some of the most influential papers in logic and the philosophy of logic, such as Tarski (1944), Turing (1950), Quine (1951), and Davidson (1967), were published in so-called 'general' philosophy journals, such as *Philosophy and Phenomenological Research*, *Mind*, *Philosophical Review*, and *Synthese*, which are included in JSTOR's Philosophy category.

Much like philosophy, mathematics is supposed to be like logic in terms of being a priori (cf. Ashton and Mizrahi 2018), which is why I have selected JSTOR's Mathematics discipline category for the sake of comparison as well. I have also selected the Biological Sciences category in JSTOR because biology is supposed to be different from logic insofar as the latter is supposed to be a priori, whereas the former is supposed to be a posteriori.

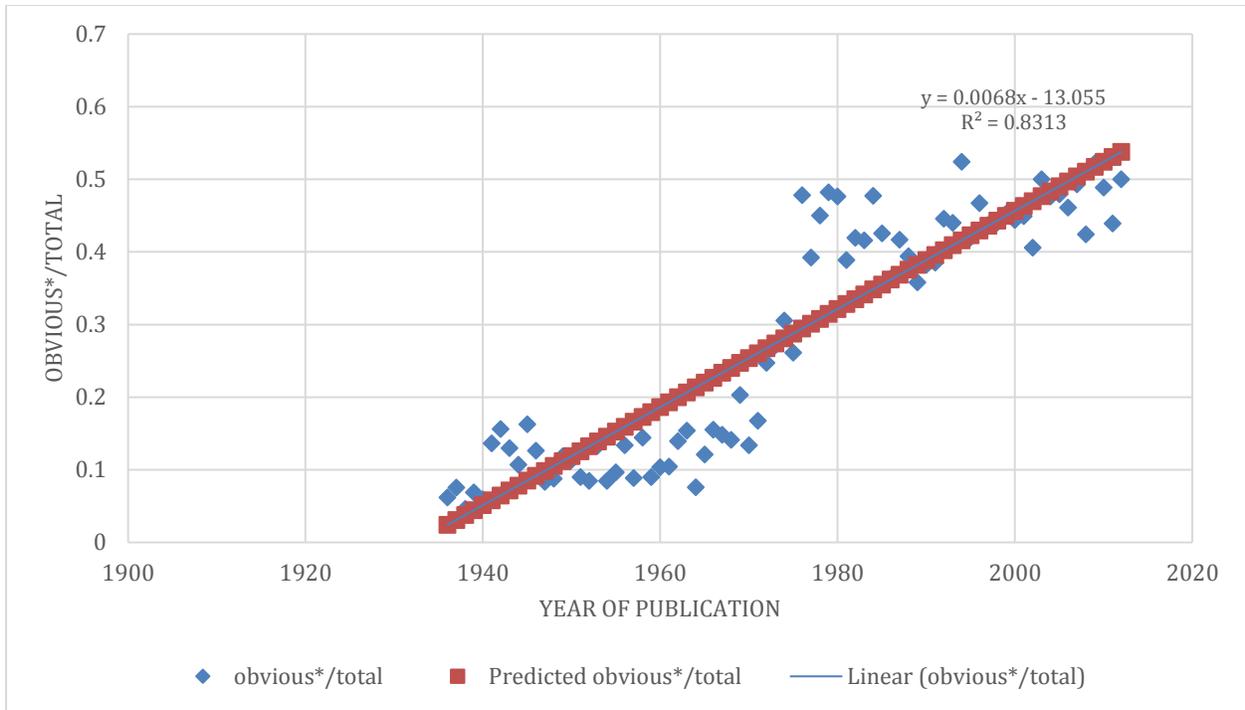
3. Results

Before testing predictions (A)-(C), it would be useful to get a sense of how widespread ‘obvious’ talk is in research articles on logic, philosophical logic, and the philosophy of logic, and to compare that to how often ‘obvious’ is used in our comparison disciplines, namely, philosophy, mathematics, and biology. So I first looked at the number of Logic, Philosophy, Mathematics, and Biology research articles that contain the word ‘obvious’ and its cognates (‘obviously’ and ‘obviousness’) in the JSTOR corpus.⁹ As mentioned above, JSTOR allows for truncation or ‘wildcard’ search, so I used the search term ‘obvious*’.

Figure 1 shows the proportions of research articles (written in English) in the Logic dataset that contain the word ‘obvious’ and its cognates (from 1936 until 2012).

Figure 1. Proportions of research articles in the Logic dataset that contain ‘obvious’ and its cognates (Source: JSTOR Data for Research)

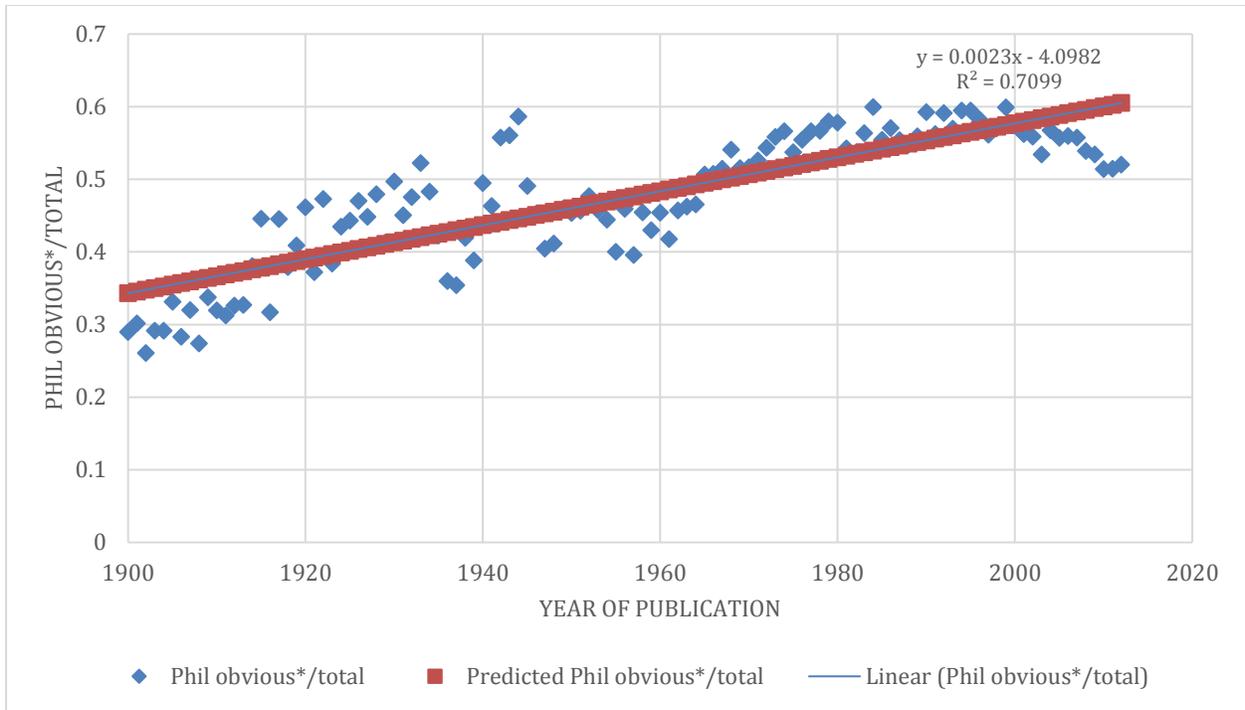
⁹ When ‘logic’ is not capitalized, it refers to the field of study or academic discipline of logic. When ‘Logic’ is capitalized, it refers to the dataset of logic articles mined from the JSTOR database. The same applies to ‘philosophy’ (field) and ‘Philosophy’ (dataset), ‘mathematics’ (field) and ‘Mathematics’ (dataset), and ‘biology’ (field) and ‘Biology’ (dataset).



As far as the Logic dataset is concerned, a linear regression analysis indicates that year predicts proportion of ‘obvious’ talk in logic research articles. A significant regression equation was found ($F(1, 75) = 369.58, p = .00$), with an adjusted R^2 of .83. This suggests that ‘obvious’ talk in logic research articles is becoming significantly more widespread over the years.

Figure 2 shows the proportions of research articles (written in English) in the Philosophy dataset that contain the word ‘obvious’ and its cognates (from 1900 until 2012).

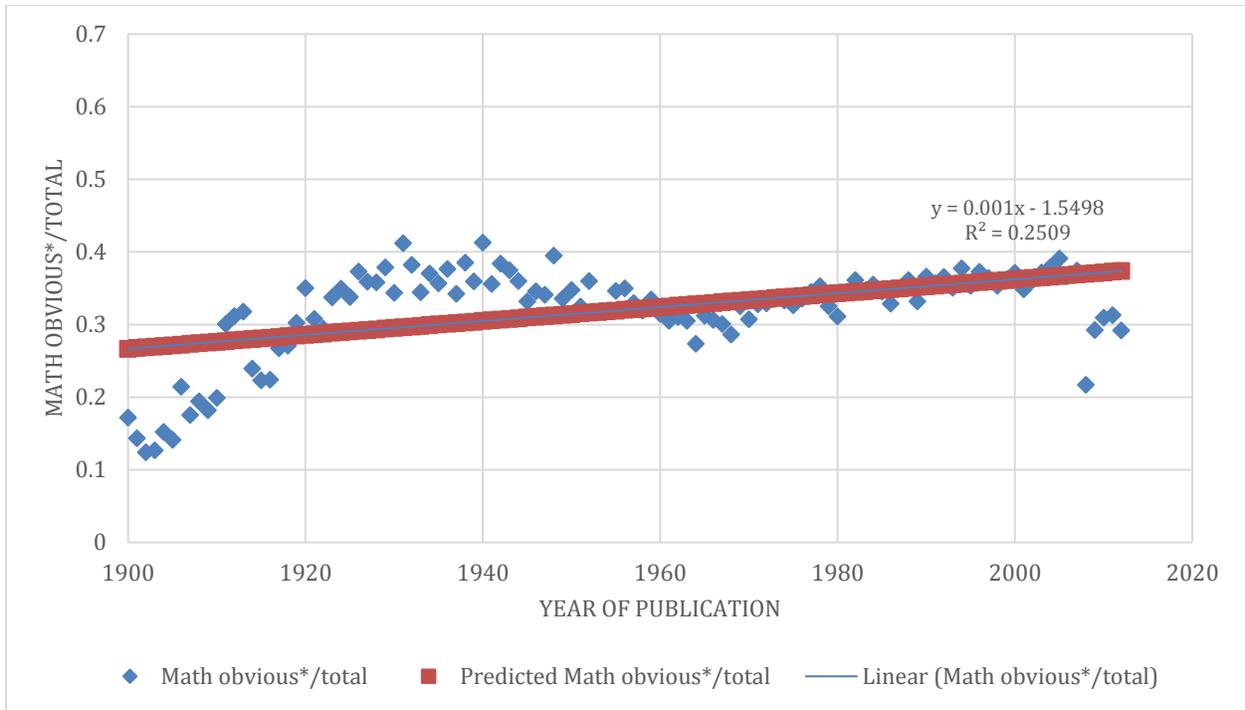
Figure 2. Proportions of research articles in the Philosophy dataset that contain ‘obvious’ and its cognates (Source: JSTOR Data for Research)



As far as the Philosophy dataset is concerned, a linear regression analysis indicates that year predicts proportion of ‘obvious’ talk in philosophy research articles. A significant regression equation was found ($F(1, 111) = 271.57, p = .00$), with an adjusted R^2 of .71. This suggests that ‘obvious’ talk in philosophy research articles is becoming significantly more widespread over the years.

Figure 3 shows the proportions of research articles (written in English) in the Mathematics dataset that contain the word ‘obvious’ and its cognates (from 1900 until 2012).

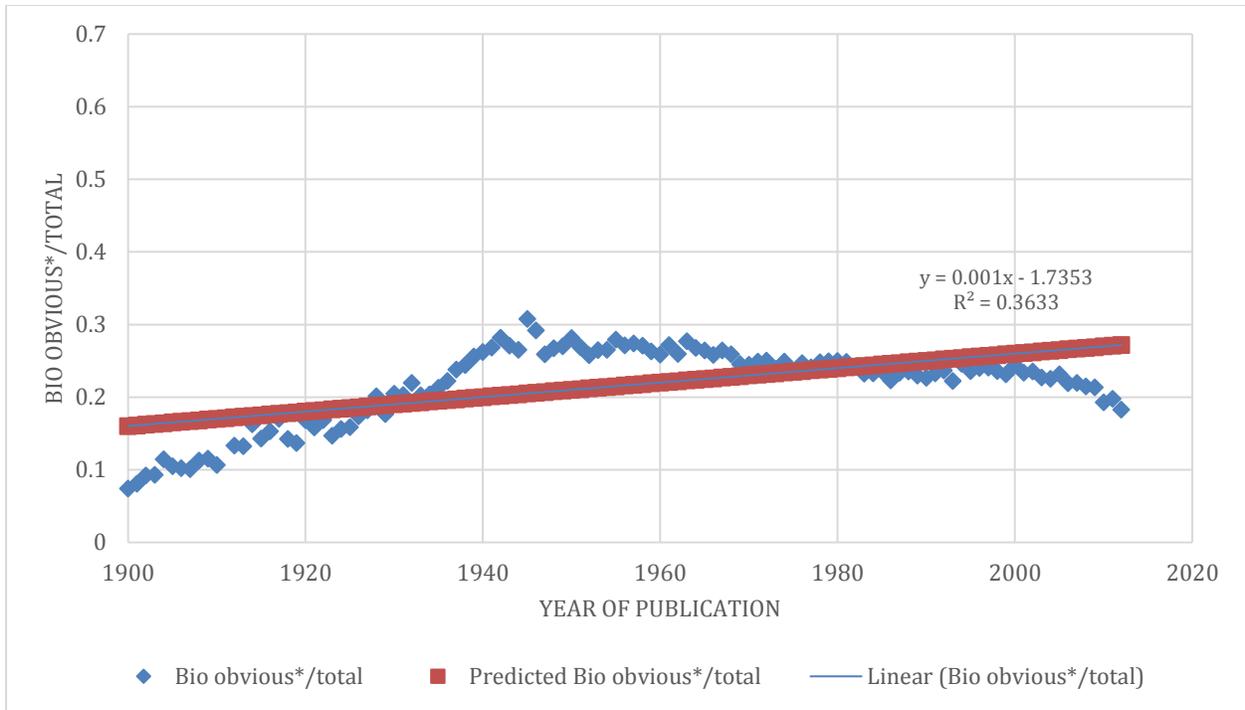
Figure 3. Proportions of research articles in the Mathematics dataset that contain ‘obvious’ and its cognates (Source: JSTOR Data for Research)



As far as the Mathematics dataset is concerned, a linear regression analysis indicates that year predicts proportion of ‘obvious’ talk in mathematics research articles. A significant regression equation was found ($F(1, 111) = 37.19, p = .00$), with an adjusted R^2 of .25. This suggests that ‘obvious’ talk in mathematics research articles is becoming significantly more widespread over the years.

Finally, figure 4 shows the proportions of research articles (written in English) in the Biology dataset that contain the word ‘obvious’ and its cognates (from 1900 until 2012).

Figure 4. Proportions of research articles in the Biology dataset that contain ‘obvious’ and its cognates (Source: JSTOR Data for Research)



As far as the Biology dataset is concerned, a linear regression analysis indicates that year predicts proportion of ‘obvious’ talk in biology research articles. A significant regression equation was found ($F(1, 111) = 63.35, p = .00$), with an adjusted R^2 of .36. This suggests that ‘obvious’ talk in biology research articles is becoming significantly more widespread over the years.

Overall, these results suggest that ‘obvious’ talk is quite widespread in logical and philosophical practice (Table 2) and that it is becoming significantly more widespread over the years. As we have seen, in Philosophy and in Logic, year explains 71% and 83% of the variation in proportion of ‘obvious’ talk in research articles, respectively, whereas in Mathematics and Biology it explains 25% and 36%, respectively.

Table 2. Mean proportions of research articles that contain the word ‘obvious’ and its cognates relative to the total number of research articles in the Logic (1936-2012), Philosophy (1900-2012), Mathematics (1900-2012), and Biology (1900-2012) datasets

	<i>Mean</i>	<i>SD</i>	<i>N</i>
Philosophy	.48	.09	112
Mathematics	.33	.06	112
Logic	.29	.16	77
Biology	.22	.05	112

Now that we have a sense of how widespread ‘obvious’ talk is in Logic, Philosophy, Mathematics, and Biology, we can turn to testing predictions (A)-(C). If practitioners indeed understand logic as obvious, necessary, and a priori, we would also expect instances of ‘obvious’ to occur in the context of deductive arguments, since unlike inductive arguments, deductive arguments are the sort of arguments that are said to be ‘indefeasible’ and whose premises ‘provide conclusive support for their conclusions’ or ‘necessitate the truth of the conclusion’ (Sinnott-Armstrong and Fogelin 2015: 181). For example, if one thinks, as Mittelstaedt and Stachow (1978: 184) do, that, in the calculus Q_{eff} of effective quantum logic, ‘It is *obvious* that, without further knowledge about the mutual commensurability of a and b , it is impossible to decide whether the proposition $a \wedge b$ is true or false’ (emphasis added), then it would be rather odd if one were to preface the conclusion that ‘it is impossible to decide whether the proposition $a \wedge b$ is true or false’ with ‘probably’, ‘likely, or ‘it seems’, for this conclusion follows necessarily from what is taken to be obvious here. In general, then, if logicians and philosophers of logic subscribe to the view that logic is obvious, necessary, and a priori, then we would expect to see them make deductive arguments more often than other kinds of argument when they do

logic or talk about logic, for deductive arguments are the sort of arguments whose conclusions follow necessarily from their premises.

To find out if this is indeed the case, I tested for correlations between ‘obvious’ and the deductive indicator pairs, inductive indicator pairs, and hedging indicator pairs listed in Table 1 in the Logic, Philosophy, Mathematics, and Biology datasets. The results are summarized in Table 3.

Table 3. Pearson correlation coefficients for ‘obvious’ and the deductive, inductive, and hedging indicator pairs in the Logic, Philosophy, Mathematics, and Biology datasets

	Indicator Pairs	Logic	Philosophy	Mathematics	Biology
Deductive	obvious* & therefore necessarily	.47	.18	.01	.25
	obvious* & therefore certainly	.49	-.21	.18	.05
	obvious* & follows necessarily	.48	.21	.46	-.23
	obvious* & follows certainly	.42	.04	.03	-.24
Inductive	obvious* & therefore probably	-.01	-.03	-.10	.63
	obvious* & therefore likely	.31	.52	.37	.31
	obvious* & follows probably	.15	.01	-.03	.01
	obvious* & follows likely	.17	.51	.24	.49
Hedging	obvious* & therefore seem	.37	-.26	-.05	.01
	obvious* & therefore suggest	.29	.55	.36	.32
	obvious* & follows seem	.43	.37	.20	-.33
	obvious* & follows suggest	.39	.55	.16	.40

3a. ‘Obvious*’ and deductive indicator pairs

As mentioned in Section 2, if practitioners understand logic as obvious, necessary, and a priori, we would expect ‘obvious’ to positively correlate with deductive indicator pairs, such as ‘therefore necessarily’ and ‘follows certainly’. In other words, if practitioners understand logic as obvious, necessary, and a priori, then we would expect them to make arguments that reflect that obviousness in the form of deductive arguments whose conclusions *follow necessarily* from their premises.

As we can see from Table 3, in the Logic dataset, ‘obvious*’ is positively correlated with the deductive indicator pairs ‘therefore necessarily’ ($r = .47$), ‘therefore certainly’ ($r = .49$), ‘follows necessarily’ ($r = .48$), and ‘follows certainly’ ($r = .42$), as we would expect if logicians and philosophers of logic subscribe to the idea that logic is obvious, necessary, and a priori. In the Philosophy dataset, however, things are rather more complicated in practice than they are in theory. For the deductive indicator pair ‘therefore certainly’ is negatively correlated with ‘obvious*’ ($r = -.21$). Given that the Pearson correlation coefficient r can tell us about the linear relationship between two variables (positive or negative) and the strength of that relationship (the closer r is to 0, the weaker the linear relationship; the closer r is to -1, the stronger the negative relationship; the closer r is to 1, the stronger the positive relationship), we can say that the correlations between ‘obvious*’ and the deductive indicator pairs ‘therefore necessarily’ ($r = .18$), ‘follows necessarily’ ($r = .21$), and ‘follows certainly’ ($r = .04$) are rather weak as far as the Philosophy dataset is concerned. Indeed, the positive correlation between the deductive indicator pair ‘therefore necessarily’ and ‘obvious*’ is stronger in Biology ($r = .25$) than in Philosophy ($r = .18$), which is contrary to what we would expect if unlike logic and philosophy, which are supposed to be a priori, biology is supposed to be a posteriori. Although mathematics is supposed to be like logic in terms of being a priori, and ‘obvious*’ is positively correlated with

deductive indicator pairs in the Mathematics dataset, these positive correlations are generally weaker than those in the Logic dataset, with the exception of ‘follows necessarily’ ($r = .46$).

If practitioners understand logic as obvious, necessary, and a priori, we would expect the correlations between ‘obvious*’ and deductive indicator pairs to be stronger in Philosophy and Logic than in Biology. Although the deductive indicator pairs ‘follows necessarily’ and ‘follows certainly’ are negatively correlated with ‘obvious*’ in Biology, the deductive indicator pair ‘therefore certainly’ is positively correlated with ‘obvious*’ in Biology, but not in Philosophy, and the correlation between the deductive indicator pair ‘therefore necessarily’ and ‘obvious*’ in Philosophy is weaker than that in Biology. This result is not quite what we would expect to find in logical and philosophical practice because, if practitioners subscribe to the idea that logic is obvious, necessary, and a priori, then logical and philosophical practice should reflect that with language closer to that of mathematical practice than biological practice. For, unlike logic, philosophy, and mathematics, which are supposed to be a priori, biology is supposed to be a posteriori. That is, deductive indicator pairs should be positively correlated with ‘obvious*’ at least as strongly as ‘follows necessarily’ is correlated with ‘obvious*’ in Mathematics ($r = .46$). In the Logic dataset, ‘obvious*’ is positively correlated with all the deductive indicator pairs and somewhat more strongly than in Mathematics.

In terms of prediction (A), then, the results do meet our expectations overall. Prediction (A) is borne out by the data to the extent that there are positive correlations between ‘obvious*’ and the deductive indicator pairs in Logic. However, there is also a negative correlation between ‘obvious*’ and a deductive indicator pair in Philosophy, and the positive correlations are rather weak, sometimes even weaker than those in a supposedly a posteriori discipline like biology in which conclusions are tentative, not obvious.

3b. 'Obvious*' and inductive indicator pairs

As mentioned in Section 2, if practitioners understand logic as obvious, necessary, and a priori, we would expect 'obvious' to not only positively correlate with deductive indicator pairs, such as 'therefore necessarily' and 'follows certainly', but also to negatively correlate (or not correlate at all) with inductive indicator pairs, such as 'therefore probably' and 'follows likely'. In other words, if practitioners understand logic as obvious, necessary, and a priori, then we would expect them to make arguments that reflect that, namely, deductive arguments whose premises necessarily entail their conclusions, not inductive arguments that provide probable, but not conclusive, support for their conclusions.

As we can see from Table 3, as far as research articles in logic, philosophical logic, and the philosophy of logic are concerned, 'obvious' talk is positively correlated with inductive indicator pairs, with the exception of 'therefore probably' ($r = -.01$ in Logic and $r = -.03$ in Philosophy), though these negative correlations are rather weak. In Mathematics, the inductive indicator pairs 'therefore probably' and 'follows probably' are negatively correlated with 'obvious*' ($r = -.10$ and $r = -.03$), as we might expect, since mathematics, like logic and philosophy, is supposed to be a priori, but 'obvious*' is positively correlated with 'therefore likely' ($r = .37$) and 'follows likely' ($r = .24$). In Philosophy, however, the inductive indicator pairs, with the exception of 'therefore probably', are positively correlated with 'obvious*'. The strength of the positive correlations between 'obvious*' and the inductive indicator pairs 'therefore likely' ($r = .52$) and 'follows likely' ($r = .51$) approximates that between 'obvious*' and inductive indicator pairs in Biology ($r = .63$ and $r = .49$), as does the strength of the positive correlation between 'obvious*' and 'therefore likely' in Logic ($r = .31$). In that respect, logical

and philosophical practice look more like biological practice than mathematical practice, which is contrary to what we would expect if logicians and philosophers of logic take logic to be obvious, necessary, and a priori, given that logic, like mathematics and philosophy, is supposed to be a priori. For, if they take logic to be obvious, necessary, and a priori, then we would expect to see logicians and philosophers of logic make the sort of arguments whose premises *necessitate* (rather than make more probable) the truth of their conclusions. In other words, we would expect to see logicians and philosophers of logic make *deductive*, rather than inductive, arguments.

If practitioners understand logic as obvious, necessary, and a priori, we would expect the correlations between ‘obvious*’ and inductive indicator pairs to be stronger in Biology than in Philosophy or Logic, since biology is supposed to be a posteriori, whereas logic and philosophy are supposed to be a priori. As we have seen, however, that is not quite what we find in practice. With the exception of ‘therefore probably’, which is negatively correlated with ‘obvious*’ in Philosophy ($r = -.03$), the other inductive indicator pairs are not only positively correlated with ‘obvious*’ in Philosophy but the strength of these correlations approximates that of positive correlations in Biology. Moreover, ‘obvious’ talk in Logic is also positively correlated with the inductive indicator pairs with strength that approximates that of a positive correlation in Biology. Contrary to expectations, these results suggest that, as far as scholarly work in logic, philosophical logic, and the philosophy of logic is concerned, ‘obvious’ talk is likely to occur not only in the context of deductive argumentation but also in the context of inductive argumentation. These results are at odds with the way in which we would expect to see the idea that logic is obvious reflected in logical and philosophical practice because the conclusions of inductive arguments are probable, not certain, and thus always tentative or doubtful to some extent, whereas obvious and necessary truths are supposed to be neither tentative nor doubtful.

In terms of prediction (B), then, the results are rather mixed. Prediction (B) is partially borne out by the data to the extent that there is a negative correlation between ‘obvious*’ and an inductive indicator pair in both Logic and Philosophy. However, there are more positive than negative correlations between ‘obvious*’ and inductive indicator pairs in both Logic and Philosophy, and those positive correlations are often as strong as those in a supposedly a posteriori discipline like biology in which conclusions are tentative, not obvious or necessary, whereas a supposedly a priori discipline like mathematics shows some negative correlations between ‘obvious*’ and inductive indicator pairs.

3c. ‘Obvious*’ and hedging indicator pairs

As mentioned in Section 2, if practitioners understand logic as obvious, necessary, and a priori, we would expect ‘obvious’ to not only positively correlate with deductive indicator pairs, such as ‘therefore necessarily’ and ‘follows certainly’, and negatively correlate (or not correlate at all) with inductive indicator pairs, such as ‘therefore probably’ and ‘follows likely’, but also to negatively correlate (or not correlate at all) with hedging indicator pairs, such as ‘therefore seem’ and ‘follows suggest’. After all, there is no need to hedge an obvious, necessary, and a priori truth. For instance, if a statement like ‘All squares are rectangles’ is obvious, necessary, and a priori, then there is no need to preface it with ‘it seems that’, as in ‘It seems that all squares are rectangles’. Likewise, if ‘Double negation is fairly obvious and needs little explanation’ (Hurley 2012: 402), then it would be rather odd to preface it with ‘evidence suggests that’, as in ‘Our evidence suggests that $p \equiv \sim \sim p$ ’.

As we can see in Table 3, all the hedging indicator pairs are positively correlated with ‘obvious*’ in the Logic dataset. These results suggest that logical practice is rather like

biological practice, given that the hedging indicator pairs ‘therefore suggest’ ($r = -.29$ in Logic and $r = .32$ in Biology) and ‘follows suggest’ ($r = -.39$ in Logic and $r = .40$ in Biology) are positively correlated with ‘obvious*’ in Biology, as one might expect, since biology, unlike logic, is supposed to be an a posteriori discipline in which conclusions are tentative, not obvious or necessary. In Philosophy, too, the hedging indicator pairs, with the exception of ‘therefore seem’ ($r = -.26$), are positively correlated with ‘obvious*’. Indeed, these correlations between hedging indicator pairs and ‘obvious*’ are generally stronger in Philosophy than in Biology. Contrary to expectations, then, these results show that hedging indicator pairs, such as ‘therefore suggest’ ($r = .29$ in Logic and $r = .55$ in Philosophy) and ‘follows seem’ ($r = .43$ in Logic and $r = .37$ in Philosophy), are positively correlated with ‘obvious*’, thus suggesting that logicians and philosophers of logic hedge their claims in argumentative contexts when they do logic or talk about the obviousness of logic.

If logicians and philosophers of logic take logic to be obvious, necessary, and a priori, we would expect the correlations between ‘obvious*’ and hedging indicator pairs to be stronger in Biology than in Logic or in Philosophy, given that biology is supposed to be a posteriori, where conclusions are tentative, whereas logic and philosophy are supposed to be a priori, where conclusions are obvious and necessary. As we have seen, however, that is not quite what we find in practice. With the exception of ‘therefore seem’, which is negatively correlated with ‘obvious*’ in Philosophy ($r = -.26$), the other hedging indicator pairs are not only positively correlated with ‘obvious*’ in Philosophy but also more strongly correlated than in Biology. In fact, in the Logic dataset, we see positive correlations between ‘obvious*’ and all the hedging indicator pairs that are generally stronger than positive correlations in Biology. These results are at odds with the way in which we would expect to see the idea that logic is obvious reflected in

logical and philosophical practice because there is no need to hedge claims that are obvious, necessary, and a priori. Again, if a statement like ‘All squares are rectangles’ is obvious, necessary, and a priori, then it would be odd to say ‘It seems that all squares are rectangles’ or ‘Our evidence suggests that all squares are rectangles’. Even though, on the idea that ‘logic is a theory of the obvious’ (Sher 1999: 207), logical truths are supposed to be like ‘All squares are rectangles’ insofar as they are supposed to be obvious, necessary, and a priori, the results from the Logic dataset suggest that logicians and philosophers of logic use hedging markers when they make such statements in argumentative contexts.

In terms of prediction (C), then, the results are somewhat mixed as well. While there is a negative correlation between ‘obvious*’ and a hedging indicator pair in Philosophy, which is supposed to be a priori like logic, most of the correlations between ‘obvious*’ and hedging indicator pairs are positive in both Logic and Philosophy. In fact, the Logic dataset shows positive correlations between ‘obvious*’ and the hedging indicator pairs that are generally stronger than positive correlations in Biology, even though biology is supposed to be different from both logic and philosophy in terms of being an a posteriori field with tentative conclusions rather than an a priori field with obvious and necessary conclusions.

4. Discussion

The results of my empirical study can be summed up as follows. As we would expect to see reflected in logical practice (i.e., in research articles on logic, philosophical logic, and philosophy of logic) if practitioners subscribe to the view that ‘logic is a theory of the obvious’ (Sher 1999: 207) and that ‘logic is [...] necessary and a priori’ (Russell 2015: 793):

1. The word ‘obvious’ (and its cognates) is positively correlated with the deductive indicator pairs ‘therefore necessarily’ ($r = .47$ in Logic and $r = .18$ in Philosophy), ‘therefore certainly’ ($r = .49$ in Logic), ‘follows necessarily’ ($r = .48$ in Logic and $r = .21$ in Philosophy), and ‘follows certainly’ ($r = .42$ in Logic and $r = .04$ in Philosophy).
2. The word ‘obvious’ (and its cognates) is negatively correlated with the inductive indicator pair ‘therefore probably’ in Logic ($r = -.01$) and in Philosophy ($r = -.03$).
3. The word ‘obvious’ (and its cognates) is negatively correlated with the hedging indicator pair ‘therefore seem’ in Philosophy ($r = -.26$).
4. Generally, the strongest positive correlations between ‘obvious’ talk and the deductive indicator pairs are found in the Logic dataset.

Contrary to what we would expect to see reflected in logical practice (i.e., in research articles on logic, philosophical logic, and philosophy of logic) if practitioners subscribe to the view that ‘logic is a theory of the obvious’ (Sher 1999: 207) and that ‘logic is [...] necessary and a priori’ (Russell 2015: 793):

5. Although positive, the correlations between the word ‘obvious’ (and its cognates) and deductive indicator pairs in Philosophy are rather weak.
6. The word ‘obvious’ (and its cognates) is negatively correlated with the deductive indicator pair ‘therefore certainly’ in Philosophy ($r = -.21$).
7. The word ‘obvious’ (and its cognates) is positively correlated with the inductive indicator pairs ‘therefore likely’ ($r = .31$ in Logic and $r = .52$ in Philosophy), ‘follows probably’ ($r = .15$ in Logic and $r = .01$ in Philosophy), and ‘follows likely’ ($r = .17$ in Logic and $r = .51$ in Philosophy).

8. The word ‘obvious’ (and its cognates) is positively correlated with the hedging indicator pairs ‘therefore seem’ ($r = .37$ in Logic), ‘therefore suggest’ ($r = .29$ in Logic and $r = .55$ in Philosophy), ‘follows seem’ ($r = .43$ in Logic and $r = .37$ in Philosophy), and ‘follows suggest’ ($r = .39$ in Logic and $r = .55$ in Philosophy).
9. Generally, the positive correlations between ‘obvious’ talk and the inductive indicator pairs are as strong in Philosophy as they are in Biology.
10. Generally, the positive correlations between the word ‘obvious’ (and its cognates) and the hedging indicator pairs are as strong in Logic as they are in Biology.

These mixed results do not warrant any definitive conclusions about the status of the idea that logic is obvious in logical practice. Indeed, these mixed results suggest that the idea that logic is obvious is not reflected in any straightforward way in the arguments that logicians and philosophers of logic actually make in their scholarly work. Instead, they suggest that things are rather more complicated in practice than they are in theory. For although many logicians and philosophers of logic subscribe to the idea ‘logic is a theory of the obvious’ (Sher 1999: 207) and that ‘logic is [...] necessary and a priori’ (Russell 2015: 793), as evidenced by the fact that ‘obvious’ talk is becoming significantly more widespread in logic (Figure 1) and philosophy (Figure 2) over the years, the arguments they actually put forward in their scholarly work appear in the context of inductive indicators and hedging markers. This is contrary to the sort of arguments we would expect to see in research articles on logic, philosophical logic, and philosophy of logic if practitioners took logic to be obvious, necessary, and a priori.

For this reason, I submit, the results of my empirical study raise an interesting question that is worthy of further investigation: if logicians and philosophers of logic subscribe to the view that logic is obvious, necessary, and a priori, why do they hedge their claims and frame

their arguments using inductive indicators? The fact that they do so in practice is *prima facie* in tension with the idea that logic is obvious, necessary, and a priori. This fact, then, calls for an explanation. Providing such an explanation is beyond the scope of this chapter, so I hope that future studies will be able to shed light on this issue.¹⁰

Beyond the aforementioned results, I think that my empirical study speaks to the broader question of how studying philosophical texts from a data science perspective can shed new light on logical, philosophical, and metaphilosophical questions. While some of the data I have mined from the JSTOR database do indeed meet our expectations, there are some surprising findings as well. One surprising finding, as mentioned above, is that ‘obvious’ talk is becoming significantly more widespread in logic (Figure 1) and philosophy (Figure 2) over the years, but that such talk also occurs in the context of inductive argumentation and hedging markers. This suggests that, although they subscribe to the view that logic is obvious, logicians and philosophers of logic engage in inductive argumentation and hedge their claims when they do logic or talk about the obviousness of logic. This finding could not have come to light without the data mining and corpus analysis techniques of data science. This, then, is an illustration of how data science can help us gain new insights about logic, philosophy, and other disciplines as well, when we apply its methods to analyzing large corpora of scholarly work in those fields of study.

My hope is that this data-driven methodology can be used to investigate systematically other claims that are often made about fields like logic and philosophy as well as other disciplines.¹¹ As more refined techniques of data mining and analysis become available, such

¹⁰ In that respect, it is interesting to point out the increase in ‘obvious’ talk around the 1970s (see Figure 1) and the fact that the *Journal of Philosophical Logic* (JPL) published its first issue in 1972. Further studies are needed to determine whether, and the extent to which, JPL is responsible for the significant increase in ‘obvious’ talk in logic as a whole. Thanks to Andrew Aberdein for this point.

¹¹ For another example of applying the methods of data science to philosophy, see Ashton and Mizrahi (2017).

that they would allow us to conduct more focused studies using specialized corpora, we will be able to gain a better understanding of the ways in which practitioners in philosophy, mathematics, logic, and biology conceive of their arguments and methods. In that respect, the application of data science to logic, philosophical logic, and the philosophy of logic can help initiate a discussion on the methodology of logic and philosophy of logic just as the application of social science to philosophy, also known as ‘experimental philosophy’, has engendered a fruitful discussion about philosophical methodology in recent years. As Dolcini (2017: 102) puts it, ‘The advent of experimental philosophy [...] revitalized the discussion about a major meta-philosophical issue: what are the proper methods, aims and ambitions of philosophy?’ Data science, I propose, can do for logic, philosophical logic, and the philosophy of logic in particular what experimental philosophy did for philosophy in general.

5. Conclusion

In this chapter, I have looked at one way in which a data science approach to studying logical and philosophical texts can shed new light on logical, philosophical, and metaphilosophical questions. In particular, I have tested empirically how the idea that logic is obvious, necessary, and a priori is reflected in logical and philosophical practice. The results of my empirical survey of data mined systematically from the JSTOR database suggest that there is a difference between theory and practice as far as the idea that logic is obvious is concerned. That is, although my results suggest that logicians and philosophers of logic subscribe to the view that logic is obvious, necessary, and a priori *in theory*, given that ‘obvious’ talk is becoming significantly more widespread in logic and philosophy over the years, they also show that, *in practice*, ‘obvious’ talk in logic and philosophy often occurs in the context of inductive argumentation and

hedging markers. This is contrary to the sort of arguments we would expect to see in research articles on logic, philosophical logic, and the philosophy of logic if practitioners understand logic to be obvious, necessary, and a priori.

These findings, I submit, raise the following interesting question for further research: if logicians and philosophers of logic subscribe to the view that logic is obvious, necessary, and a priori, as evidenced by the fact that ‘obvious’ talk is becoming significantly more widespread in both logic and philosophy over the years, why do they hedge their claims and frame their arguments using inductive indicators when they use the term ‘obvious’ and its cognates? I have proposed that addressing such questions using the methods of data science can help initiate a discussion on methodology in logic and the philosophy of logic in much the same way that addressing philosophical questions using the methods of social science (i.e., ‘experimental philosophy’) has led to a fruitful discussion about philosophical methodology.

Acknowledgments

I would like to thank Andrew Aberdein and Matthew Inglis for inviting me to contribute to this volume and to two anonymous reviewers for helpful comments on an earlier draft of this chapter.

References

Ashton, Z. and M. Mizrahi (2017), ‘Intuition Talk is Not Methodologically Cheap: Testing the “Received Wisdom” about Armchair Philosophy’, *Erkenntnis*, <https://doi.org/10.1007/s10670-017-9904-4>.

Ashton, Z. and M. Mizrahi (2018), 'Show Me that Argument: Empirically Testing the Armchair Philosophy Picture', *Metaphilosophy*, 49 (1-2): 58-70.

Baronett, S. (2016), *Logic*, New York: Oxford University Press.

Bessie, J. and S. Glennan (2000), *Elements of Deductive Inference*, Belmont, CA: Wadsworth.

Bueno, O. (2010), 'Is Logic A Priori?' *The Harvard Review of Philosophy*, 17 (1): 105-117.

Cahn, S. M. (2017), 'How Teachers Succeed', *Blog of the APA*, 11 May. Available online: <http://blog.apaonline.org/2017/05/11/how-teachers-succeed/>.

Cappelen, H. (2012), *Philosophy without Intuitions*, Oxford: Oxford University Press.

Copi, I. M., C. Cohen and K. McMahon (2011), *Introduction to Logic*, 14th edn, New York: Prentice Hall.

Davidson, D. (1967). 'Truth and Meaning', *Synthese*, 17 (3): 304-323.

De Morgan, A. (1839). *First Notions of Logic (Preparatory to the Study of Geometry)*. London: Printed for Taylor and Walton, Booksellers and Publishers to University College.

Dolcini, N. (2017). 'Philosophy Made Visual: An Experimental Study', in L. Magnani and C. Casadio (eds.), *Model-Based Reasoning in Science and Technology: Logical, Epistemological, and Cognitive Issues*, 101-118, Basel: Springer.

Field, H. (1998), 'Epistemological Nonfactualism and the A Prioricity of Logic', *Philosophical Studies*, 92 (1): 1-24.

Field, H. (2009), 'What Is the Normative Role of Logic?' *The Aristotelian Society Supplementary Volume*, 83 (1): 251-268.

Horsten, L. (2015), 'One Hundred Years of Semantic Paradox', *Journal of Philosophical Logic*, 44 (6): 681-695.

Hurley, P. J. (2012), *A Concise Introduction to Logic*, 11th edn, Boston, MA: Wadsworth.

Jenkins, K. (2014). "'That's Not Philosophy": Feminism, Academia, and the Double Bind', *Journal of Gender Studies*, 23 (3): 262-274.

Kauppinen, A. (2007). 'The Rise and Fall of Experimental Philosophy', *Philosophical Explorations*, 10 (2): 95-118.

Knobe, J. (2007). 'Experimental Philosophy and Philosophical Significance', *Philosophical Explorations*, 10 (2): 119-121.

Kovač, S. (2012), 'Logical Opposition and Collective Decisions', in J. Béziau and D. Jacquette (eds.), *Around and Beyond the Square of Opposition*, 341-356, Basel: Springer.

Kracht, M. (2013), 'Are Logical Languages Compositional?' *Studia Logica* 101 (6): 1319-1340.

Marcus, R. (2018). *Introduction to Formal Logic with Philosophical Applications*. New York: Oxford University Press.

Mittelstaedt, P. and Stachow, E. W. (1978). 'The Principle of Excluded Middle in Quantum Logic', *Journal of Philosophical Logic*, 7 (1): 181-208.

Murawski, R. (2014), *The Philosophy of Mathematics and Logic in the 1920s and 1930s in Poland*, Translated from Polish by Maria Kantor, Basel: Springer.

O'Neill, E. and E. Machery. (2014). 'Experimental Philosophy: What Is It Good For?' in E. Machery and E. O'Neill (eds.), *Current Controversies in Experimental Philosophy*, vii-xxix, New York: Routledge.

Petersen, U. (2000). 'Logic without Contraction as Based on Inclusion and Unrestricted Abstraction', *Studia Logica* 64 (3): 365-403.

Quine, W. V. (1951), 'Two Dogmas of Empiricism', *Philosophical Review*, 60 (1): 20-43.

Quine, W. V. (1970), *Philosophy of Logic*, Englewood Cliffs, NJ: Prentice Hall.

Quine, W. V. (1976), *The Ways of Paradox and Other Essays*, Revised and Enlarged Edition, Cambridge, MA: Harvard University Press.

Russell, G. (2015). 'The Justification of the Basic Laws of Logic', *Journal of Philosophical Logic*, 44 (6): 793-803.

Salmon, M. H. (2013), *Introduction to Logic and Critical Thinking*, sixth edn, Boston, MA: Wadsworth.

Shenefelt, M. and H. White (2013), *If A, then B: How the World Discovered Logic*, New York: Columbia University Press.

Sher, G. (1999), 'Is Logic a Theory of the Obvious?' *European Review of Philosophy*, 4: 207-238.

Sher, G. (2010). 'Epistemic Friction: Reflections on Knowledge, Truth, and Logic', *Erkenntnis*, 72 (2): 151-176.

Sinnott-Armstrong, W. and R. J. Fogelin (2015), *Understanding Arguments*, Ninth edn, Stamford, CT: Cengage Learning.

Spreen, D. (1998), 'On Effective Topological Spaces', *The Journal of Symbolic Logic*, 63 (1): 185-221.

Suppes, P. (1999). *Introduction to Logic*. Mineola, NY: Dover Publications, Inc.

Tarski, A. (1944), 'The Semantic Conception of Truth: and the Foundations of Semantics', *Philosophy and Phenomenological Research*, 4 (3): 341-376.

Thabet, R. A. (2018), 'A Cross-Cultural Corpus Study of the Use of Hedging Markers and Dogmatism in Postgraduate Writing of Native and Non-native Speakers of English', in K. Shaalan, A. E. Hassanien, and F. Tolba (eds.), *Intelligent Natural Language Processing: Trends and Applications*, 677-710, Gewerbestrasse: Springer.

Turing, A. M. (1950), 'Computing Machinery and Intelligence', *Mind*, 59 (236): 433-460.