
EMPIRICAL STUDIES ON *TRUTH* AND THE PROJECT OF RE-ENGINEERING *TRUTH*

BY

KEVIN REUTER, AND GEORG BRUN

Abstract: Most philosophers have largely downplayed any relevance of multiple meanings of the folk concept of truth in the empirical domain. However, confusions about what truth is have surged in political and everyday discourse. In order to resolve these confusions, we argue that we need a more accurate picture of how the term ‘true’ is in fact used. Our experimental studies reveal that the use of ‘true’ shows substantial variance within the empirical domain, indicating that ‘true’ is ambiguous between a correspondence and a coherence reading. We then explore the consequences of these results for the project of re-engineering *truth*.

1. *Introduction*

Since populism gathered momentum, the term ‘true’ seems to be given away so freely that the label ‘post-truth’ has become a commonplace diagnosis.¹ Of course, many people try to put up resistance to this tendency, for example, by building fact checker websites. But it appears that attempts at curtailing the inflationary use of the term ‘true’ have largely failed. Many others just continue to call statements ‘true’ sometimes even after acknowledging that they do not correspond to facts in the world. A common reaction is to call them liars (who intentionally tell falsehoods), bullshitters (who do not care about the truth), or fools (who are too stupid to understand what

This work is fully collaborative.

Pacific Philosophical Quarterly 103 (2022) 493–517 DOI: 10.1111/papq.12370

© 2021 The Authors

Pacific Philosophical Quarterly published by University of Southern California and John Wiley & Sons Ltd. This is an open access article under the terms of the Creative Commons Attribution-NonCommercial License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited and is not used for commercial purposes.

they are talking about). However, we believe these reactions – although certainly appropriate in many cases – easily miss an important point.

Fact checkers do a great service to society: facts indeed matter *a lot*. But, as we will argue in this paper, *truth* is often, especially in political and everyday discourse, taken to be not just about facts, or correspondence with reality. Other aspects – most importantly, coherence – play a crucial role as well. Although this would be reason to think that ‘true’ may be ambiguous in political and everyday discourse,² the importance of coherence can be overlooked if one is focused on statements that fall in the domain of the empirical sciences, which are committed to a notion of truth for which correspondence to reality is decisive.

Nonetheless, the idea that ‘true’ is ambiguous has long been with us. It is, for example, clearly expressed in Tarski’s groundbreaking work on truth:

The word ‘true,’ like other words from our everyday language, is certainly not unambiguous. [...] We should reconcile ourselves with the fact that we are confronted, not with one concept, but with several different concepts which are denoted by one word. (Tarski, 1944, pp. 342, 355)

Tarski links this diagnosis of ambiguity to the various philosophical conceptions of truth. Indeed, an important rival of the correspondence theory of truth is the coherence theory which holds that a statement is true if and only if it coheres with a specific set of other beliefs. Many seem to agree that in the mathematical and ethical domain, the coherence theory seems well suited to explain our practice of assigning truth.

In what follows, we present experimental studies which challenge the idea that – within the empirical domain – the common-sense notion of truth is rooted exclusively in correspondence. In these experiments, participants read vignettes in which a person makes a statement that either corresponds with reality but is incoherent with other relevant beliefs or fails to correspond with reality but is coherent with other beliefs. Perhaps surprisingly or even shockingly – at least from a philosopher’s perspective – a substantial number of participants (in some experiments up to 70%) responded in line with the predictions of the coherence account. These results suggest that, even within the empirical domain, ‘true’ is not used in a uniform way in everyday discourse. Although it has been known since Næss’s first empirical investigations and his exchange with Tarski (Næss, 1938; Tarski, 1944) that such a result might be plausible, it has neither been corroborated in systematic empirical studies nor has it been taken serious enough by philosophers.

What is more, the ambiguity of ‘true’ in everyday discourse is not harmless, given the central role this notion plays in public and political discourse – as witnessed by the recent ‘post-truth’ debate. Vital questions therefore need to be asked about the way we want our concept of truth to be. In other words, we need to determine whether and how *truth* needs to be re-engineered.³ So far, this question has been discussed in the context of

projects that aim at developing a formally rigorous account of truth (Tarski, 1933, 1944) and at dealing with liar-style paradoxes (Scharp, 2013). In contrast to these approaches, we focus on the use of ‘true’ in everyday and public discourse. As we will argue, our empirical results give us strong reasons to advocate re-engineering *truth* for these discourses. In no way, this means that we would want to advocate gerrymandering the use of ‘true’ in a way that could be seen as a legitimation for propagating coherent misinformation, whitewashing lies and bullshit, or dismissing facts and fact checking. Rather, we think that problematic aspects of using ‘true’ can and need be addressed on the basis of a more accurate picture of the everyday usage of ‘true’. More specifically, we will argue that everyday and public discourse could be improved by new ways of deploying distinctions and theoretical concepts philosophers have developed, that re-engineering *truth* requires to employ at least two target concepts, and that work on coherence may be much more pertinent than usually recognized.

In the next section, we briefly discuss how the debate about theories of truth (Section 1) and previous empirical studies (Section 2) gives us reason for thinking that ‘true’ may be ambiguous. Section 3 presents our experimental studies which indicate that ‘true’ may be ambiguous. We also discuss the most pressing challenge, which seeks to explain away our findings with the allegation that our studies do not properly distinguish *truth* from related concepts like *truthful* and *true for*. In Section 4, we explore the consequences of these results for the project of re-engineering *truth*.

2. *Background to the potential ambiguity of ‘true’*

The idea that ‘true’ may be ambiguous has been advanced from two perspectives. In Section 2.1, we briefly review how the debate about correspondence, coherence, and pluralistic theories of truth speaks in favor of thinking that ‘true’ is ambiguous. Section 2.2 summarizes some previous empirical studies, which also point in the direction of ambiguity.

2.1. CORRESPONDENCE, COHERENCE, AND PLURALISTIC THEORIES OF *TRUTH*

At the core of the many versions of the correspondence theory is the idea that only those statements are true whose content corresponds to reality: if reality is such that the Earth is getting warmer, then ‘the Earth is heating up’ is true; otherwise, it is false. The fact checkers mentioned above are clearly within the camp of the correspondence theorists. And so are many important historical figures since Greek antiquity, as well as about half of the philosophers living today (see Bourget & Chalmers, 2014). Some version of the correspondence theory is almost undoubtedly in place in the natural

sciences. It is, of course, possible that the notion of truth that the sciences employ diverges to some extent from the folk concept of truth. But without evidence to the contrary, one might assume that striving for the truth is the same endeavor whether you are engaging in a scientific exercise or in everyday talk. This assumption is further supported by claims regarding the intuitive correctness of the correspondence account. However, such intuitions lose some of their force when we move into other domains of discourse, like mathematics or ethics, which brings us to the coherence account.

The coherence theory (e.g., Putnam, 1981; Young, 2001) claims that a statement is true if and only if it coheres with a specific set of other statements or beliefs. Accordingly, a statement such as ‘snow is white’ is true, not because reality matches that content (although, of course, it still might do so), but rather because it coheres with other statements like ‘On pictures snow looks white’, ‘The reflective properties of snow are such-and-so’, and ‘If you put on tinted glasses, snow looks colored’. The coherence theory seems to be an epistemically plausible position: when people are interested in verifying or falsifying a certain belief, they will naturally consider whether other beliefs cohere or rather stand in tension with it. One might object though that we shall not confuse our epistemic practice of *finding out the truth* with what *makes true* a certain statement. While this objection seems to have some force in the empirical domain, it is less convincing in the mathematical or ethical domain. Not only do we *find out the truth* of ‘11 is a prime number’ by checking its coherence with other theorems in arithmetic, it seems to many that this statement is *made true* by cohering with fundamental mathematical axioms. Thus, at least for some areas of discourse, intuition seems to be on the coherentist side.

Not all philosophers argue for a monistic theory of truth. Pluralists about truth hold that there is more than one way for a statement to be true. Given the intuitions people have about truth in the various domains, it is not surprising that most, if not all, pluralists are scope pluralists. Thus, Edwards states that ‘the basic idea behind all forms of truth pluralism is that the analysis of truth may require different treatments for different kinds of subject matter’ (Edwards, 2012). Correspondence theories seem more convincing when applied to truths in the empirical domain. For example, ‘Water is H₂O’, and ‘Maria has got a Rolex watch in her safe’ – statements that seem to be true in virtue of corresponding to reality. Coherence theories seem intuitively plausible when accounting for truths in the legal, the ethical, or the mathematical domain. For example, ‘11 is a prime number’, and ‘Manslaughter is a felony’ – statements that seem to be true in virtue of cohering with a mathematical system, a body of law, or an ethical code.

There seem to be three ways to resolve the debate between correspondentists and coherentists – violently neglecting, of course, other accounts of truth. First, one can argue that only one of the two accounts captures the meaning of the term ‘true’. Second, one can accept that ‘true’ is

ambiguous and can express correspondence with reality, but also coherence with a certain set of other beliefs. And third, some scope pluralists have tried to cash out the meaning of ‘true’ by a set of basic platitudes that can account for the seemingly different uses in the various domains of discourse (see Wright, 2005, and Lynch, 2012). Crucially, one important factor in helping to settle this debate has been largely neglected so far. While many scholars concede that people’s intuitions about truth can be in line either with the correspondence account or with the coherence account, we do not know so far whether the concept of truth actually is ambiguous between different readings in the empirical domain of discourse.

We have so far seen that both the everyday use of the term ‘true’ and philosophical discussions on truth point towards a possible ambiguity of the folk concept of truth. Before we present our own experiments, we would like to draw on some previous empirical data that provide further reasons to take the idea that truth is ambiguous even more seriously.

2.2. PREVIOUS EMPIRICAL STUDIES ON *TRUTH*

The history of empirical studies on the common-sense notion of *truth* is itself quite remarkable and has been discussed recently in Barnard and Ulatowski (2016). In the 1930s, Arne Næss (1938) conducted interviews with around 300 laypeople. He asked his subjects explicit questions like ‘what is the common characteristic of what is true?’ Based on the highly variegated responses, he claimed that empirical investigations reveal a plurality of common-sense notions of truth, none of which dominates laypeople’s thinking. While Næss was certainly ahead of his time, asking laypeople to articulate their views about the concept of *truth* is highly problematic. For many concepts that we possess, we are capable of using them appropriately and without great effort. In contrast, correctly articulating the content of those concepts is difficult and often leads to incorrect or confabulated responses.

Apart from Næss’s later work (e.g., Næss, 1953a, 1953b), empirical investigations on the folk concept of truth were dormant for many decades. Only very recently did Barnard, Ulatowski, and Kölbel use empirical means to examine various aspects of our everyday concept of truth. In his 2008 paper ‘*True as Ambiguous*’, Kölbel (2008) considers the two statements ‘Ali G is very funny’ and ‘Statements concerning what is funny can’t be true or false’ and claims that it is quite plausible to consider both statements to be true. Assenting to both statements need not be contradictory, as long as the term ‘true’ is interpreted deflationary in the first sentence and substantially in the second. When, at the beginning of three of his philosophy courses, Kölbel asked his students to fill-in a questionnaire, a substantial amount – one time even a majority – considered ‘Ali G is funny’ to be true or false, but at the same time also considered the sentence ‘Statements concerning what is

funny cannot be true or false' to be true. Unfortunately, Kölbel does not present any numbers or statistics of the results of the survey.

Barnard and Ulatowski have advanced our knowledge of the folk concept of truth in a methodologically far more rigorous manner. They recruited hundreds of lay people on the online platform Amazon Mechanical Turk and asked them various questions to examine the way the folk understand *truth*. The collected data were subsequently subjected to statistical analysis. In Barnard and Ulatowski (2013), participants were randomly assigned to two different conditions. In the Anna condition, subjects read the sentence 'Anna has performed a simple calculation and discovered that $30 + 55 = 85$.' In the Bruno condition, participants were presented with the sentence 'Bruno has just finished painting his house. Bruno painted his house the same color as the sky on a clear summer day. Bruno claims his house is blue.' The participants were then asked to rate their agreement with the statement 'If a claim reports how the world is, then it is true' on a scale from 1 ('strongly disagree') to 5 ('strongly agree'). The results suggest that when thinking about statements in different domains of discourse – mathematical versus empirical – people will entertain a correspondence notion of truth more strongly in the empirical domain ($M_{av} = 3.63$) compared with the mathematical domain ($M_{av} = 2.81$).⁴

Both approaches certainly have their virtues. We think it is best to follow Kölbel's approach in directly asking people whether they believe a certain statement to be true or false. However, in order to examine more complex questions like whether people's concept of truth is more strongly influenced by considerations of correspondence or coherence, we need to manipulate those very aspects inside the vignettes and collect data on those responses in the rigor exemplified by Barnard and Ulatowski. The design and the results of those studies will now be presented in the next section.

3. *Experimental studies*

Both advocates of the correspondence theory and the coherence theory make falsifiable hypotheses about the common-sense notion of truth. We do believe that the hypotheses we lay out below capture an important – perhaps the most important – aspect of both theories. However, we do not commit ourselves to the implausible claim that all proponents of either of the two theories will adhere to one of those hypotheses. These are as follows:

Hypothesis 1 Correspondentism: Laypeople consider an empirical statement *S* of Person *P* to be true or false depending on whether *S* corresponds with reality, and regardless of *S*'s coherence with other relevant beliefs of Person *P*.

Hypothesis 2 Coherentism: Laypeople consider an empirical statement *S* of Person *P* to be true or false depending on whether *S* coheres with other relevant beliefs of Person *P*, and regardless of *S*'s correspondence with reality.

We merely use the terms 'correspondentism' and 'coherentism' instead of 'correspondence theory' and 'coherence theory', to indicate the somewhat restricted impact of our studies on theories of truth. For example, different versions of the coherence theory require different sets of beliefs that need to be coherent with a certain statement. While some have argued that the critical set of beliefs are those upheld by an omniscient being, others have settled for sets that encompass beliefs held by real people (Young, 2001). The coherentism hypothesis 'merely' requires coherence with other relevant beliefs of the person making the statement. We have chosen hypotheses that try to encode crucial aspects of the correspondence and coherence theories, but at the same time, allow for relatively straight-forward empirical falsification.⁵ Now, if we design vignettes that pitch coherence against non-correspondence on the one hand, and incoherence with correspondence on the other, we can investigate empirically whether one of those two hypotheses about the folk concept of truth is supported by the data. Alternatively, it might turn out that the concept of truth is ambiguous between a coherentist and a correspondentist meaning. Hence, both meanings may play a vital role in determining the truth of a statement.

We now present the methods and results of our main studies in Sections 3.1 and 3.2, before we address possible concerns and objections to our studies in Section 3.3.

3.1. STUDY 1: COHERENCE WITHOUT CORRESPONDENCE

3.1.1. *Methods*

For Study 1, 100 participants were recruited on Amazon Mechanical Turk and paid a small fee for their participation. Nine participants were excluded for indicating that they were not native English speakers. Of the remaining 91 participants, there were 50 women (and none indicating non-binary gender) with an average age of $M_{age} = 36.32$ ($SD = 11.84$). Two different vignettes were designed to avoid an influence of the scenario as an artifact.

[Party] Anne and Robert go to a party late at night. On their way to the party, Anne asks Robert whether any of his friends are at the party. Robert answers that Jill is at the party, because Jill had told Robert a few hours before that she would go. When they arrive at the party, it turns out that Jill had changed her plans, and actually is not at the party.

[Rolex] Maria is a watch collector. She keeps all her watches in a safe and knows her collection really well. One day, her friend John asks her, whether she has a 1990 Rolex Submariner in her

safe and, if so, could show it to her. Maria answers that she has got a 1990 Rolex Submariner in her safe. After all, she had purchased that watch a few years ago. When Maria opens the safe a little later, she finds out that a burglar has stolen several watches, among them the 1990 Rolex Submariner. (Note: This scenario is a variation of a vignette used by Reuter & Brössel, 2018, as well as Turri, 2013.)

In both scenarios, the main protagonist (Robert/Maria) gives an answer that is coherent with other relevant beliefs at the time in which they are asked the question. While Robert states that Jill is at the party because he had been told by her (and he had no indication that she would change her mind), Maria remembers correctly that she had purchased the Rolex watch and did not know that someone had stolen it. At the same time, both answers do not correspond to the way the world is at the time of their answer. Thus, both scenarios encode a situation in which the answers are coherent with other relevant beliefs but do not correspond with reality.

Each participant was randomly assigned to one of the two scenarios. After reading the vignette, they were asked the following question:

[Party] Was Robert's answer true or false?

[Rolex] Was Maria's answer true or false?

People were presented with three options: (1) true; (2) false; and (3) not sure.⁶

Whereas a 'true' response suggests that the participant endorses coherentism (at least for the scenario in question), a 'false' response is indicative of correspondentism. This only holds, of course, if our design really measures people's views on truth, and not related aspects like truthfulness or truth-for-ness. We will deal with those objections in Section 3.3.

3.1.2. Results

As Figure 1 shows, a majority of participants in both scenarios opted for the 'true' response, 59.6% in the Party case, and 56.8% in the Rolex case. A minority chose the 'false' response, and eight participants were 'not sure'. Pearson's χ^2 tests were used to investigate whether the results differed from the 50% mark (leaving out the 'not sure' responses). If either the correspondentism hypothesis or the coherentism hypothesis was true (see above), we would (at a minimum) expect the collected data to reveal a significant deviation from the 50% mark towards a majority of one of the two response options. However, neither in the Party case, $\chi^2 = 1.10$, $p = 0.294$, nor in the Rolex case, $\chi^2 = 2.39$, $p = 0.122$, were the results significantly different from the 50% mark.

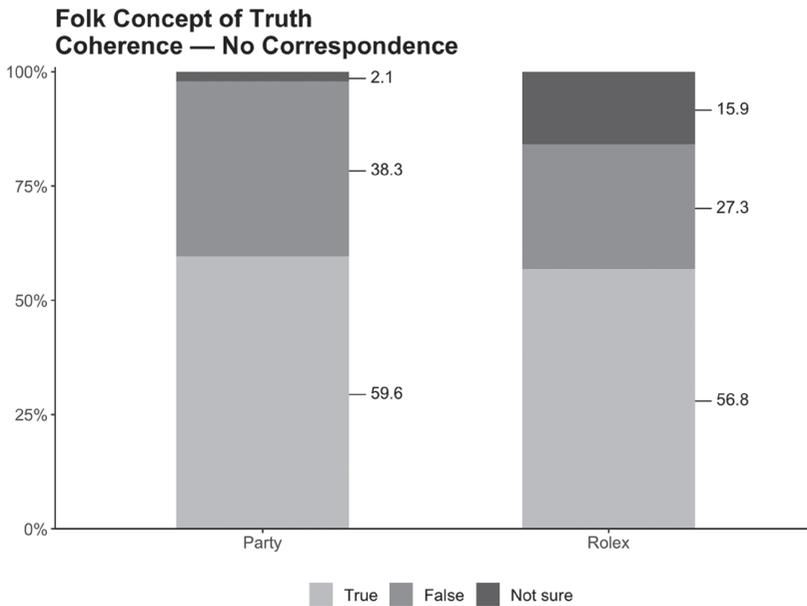


Figure 1. Responses in % to the scenarios of Study 1.

3.1.3. Discussion and follow-up study

Study 1 indicates that neither Hypothesis 1 (correspondentism) nor Hypothesis 2 (coherentism) is well supported by the empirical data we collected. Instead, some laypeople seem to be largely guided by the coherence of Maria's/Robert's answers with other relevant beliefs they hold. Other people are likely to take the lack of correspondence with reality of Maria's/Robert's answer to determine the falsity of their answer. It is also worth noting that only very few participants opted for the 'not sure' answer, suggesting that the vignette did not present a borderline case, which tells against the vagueness of 'true' in this case.

It should be noted that our study *only* investigated whether people consider Robert's or Maria's answer to be true. We have not inquired whether people are also willing to assent to a statement like 'it is true that Jill is at the party', independent of whether this statement was made by Robert. We have refrained from doing so, because asking people whether they agree with such a statement makes it no longer possible to pitch the coherentist against the correspondentist prediction. When reading the story, the reader is informed that Jill did not go to the party. Thus, the statement 'Jill is at the party' is only coherent with Robert's set of beliefs at the time in which he makes the statement, but not coherent with the reader's set of beliefs at the time in which she reads the story. On the assumption that readers will assess

'it is true that Jill is at the party' in relation to their own beliefs, the coherentist prediction therefore coincides with the correspondentist prediction.⁷

One might object, however, that we should have spelled out the propositional content that Jill is at the party, or that the Rolex watch is in the safe, when we asked participants whether Maria's/Robert's answer was true or false. We have, therefore, conducted a follow-up experiment in which we reran the experiment described above. The only difference was that we specified the propositional content of the answer. Thus, we asked the participants 'Was Robert's answer that Jill is at the party, true or false?', and, respectively, 'Was Maria's answer that she has got a 1990 Rolex Submariner in her safe, true or false?' A total of 100 participants (all English native speakers; 68 female, 32 male, and 0 non-binary; $M_{age} = 33.55$, $SD = 10.33$) were recruited on Prolific and randomly assigned to either the Party vignette or the Rolex vignette. The results were very similar to the original study. In the Party case, 23 participants selected 'true', 24 'false', and 3 'not sure'. In the Rolex case, 32 participants selected 'true', 16 'false', and 2 'not sure'. The results were not significantly different from the 50% mark (excluding 'not sure' responses) for either the Party case ($\chi^2 = 0.04$, $p = 0.837$) or the Rolex case ($\chi^2 = 2.74$, $p = 0.098$). In other words, neither of the two hypotheses is supported by the empirical data when the propositional content is made explicit in the truth question.

3.2. STUDY 2: CORRESPONDENCE WITHOUT COHERENCE

The purpose of Study 2 was to reverse the context of the scenarios of Study 1 in regard to coherence and correspondence. In Study 1, the protagonists answer coherently with their set of beliefs, but their answers do not correspond with the way the world turns out to be. In contrast, in Study 2, we designed two scenarios that pitted incoherence against correspondence.

3.2.1. *Methods*

We randomly assigned 100 participants that were recruited on Amazon Mechanical Turk to two different scenarios. Six participants were excluded for not having finished the survey or because they were not native English speakers. Of the remaining 94 participants, there were 50 women (and none indicating non-binary gender) with an average age of $M_{age} = 35.73$ ($SD = 12.93$). The two scenarios read:

[Party] Anne and Robert go to a party late at night. On their way to the party, Anne asks Robert whether any of his friends are at the party. Robert answers that Jill is at the party, although Robert had been told by Jill a few hours before that she would not go – a piece of information

that Robert completely forgot in that moment. When they arrive at the party, it turns out that Jill had changed her plans, and actually is at the party.⁸

[Rolex] Maria is a watch collector. She keeps all her watches in a safe and knows her collection really well. One day, her friend John asks her, whether she has a 1990 Rolex Submariner in her safe and, if so, could show it to her. Maria answers that she has not got a 1990 Rolex Submariner in her safe, despite the fact that Maria purchased that watch a few years ago – a piece of information that Maria completely forgot in that moment. When Maria opens the safe a little later, she finds out that a burglar has stolen several watches, among them the 1990 Rolex Submariner.

Each participant was presented with only a single vignette. After each scenario, we asked the participants to answer the same question as in Study 1.

[Party] Was Robert's answer true or false?

[Rolex] Was Maria's answer true or false?

People were presented with three options: (1) true; (2) false; and (3) not sure. This time, a 'true' response would suggest that the participant endorses correspondentism, whereas a 'false' response is indicative of coherentism.

3.2.2. *Results and discussion*

The percentages of the responses are displayed in Figure 2. In Study 2, the scenarios we used had an effect on people's responses to the questions raised in the vignettes. In the Party case, 65.2% claimed that Robert's answer was true, and 23.9% that the answer was false. In the Rolex case, the pattern was almost reversed: 54.2% decided that Maria's answer was false, and 35.4% that Maria's answer was true. Pearson's χ^2 tests revealed that the two scenarios were significantly different, both when the 'not sure' responses were included ($\chi^2 = 9.64, p = 0.008$) and when they were excluded ($\chi^2 = 9.64, p = 0.002$). We ran Pearson's χ^2 tests to find out whether the results for 'yes' and 'no' responses differed from the 50% mark (again, leaving out the 'not sure' responses). In the Party case, $\chi^2 = 4.20, p = 0.040$, the outcome was significantly different from the 50% mark. This was, however, not the case in the Rolex scenario, $\chi^2 = 0.75, p = 0.385$.

While the data in Study 1 suggest a slight majority for coherentist answers, the data are more complex to interpret in Study 2. Whereas in the Party case, more people were inclined to give a response in line with the correspondentist picture, a majority of the participants in the Rolex case were favoring the coherentist picture. However, the responses are also similar to Study 1 in suggesting that neither of the two hypotheses – correspondentism and coherentism – is supported by one of the two cases: while there was a significant majority of correspondentist answers in the

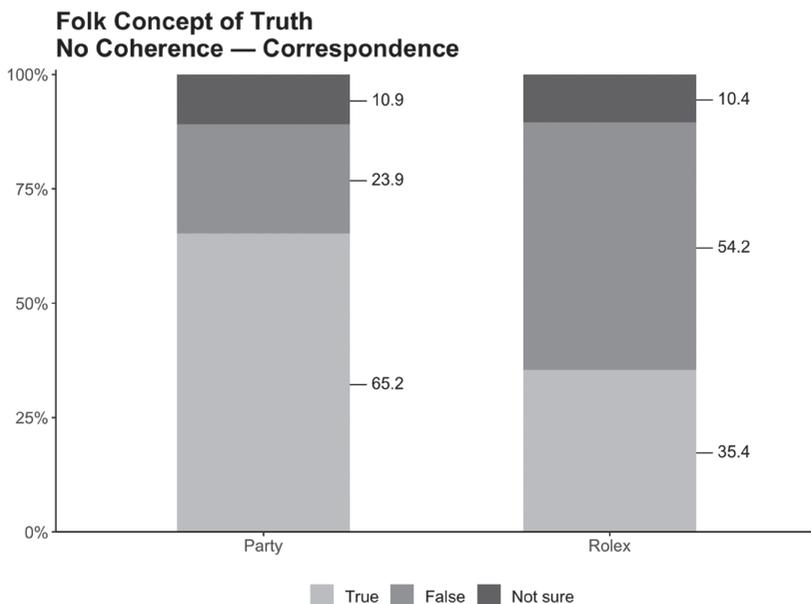


Figure 2. Responses in % to the scenarios of Study 2.

Party case, the Rolex case did not confirm this finding but rather reversed the overall pattern of results. The variance in people's responses indicates instead that the term 'true' might be ambiguous: while some people entertain a correspondentist reading of 'true', others hold a coherentist conception, at least in the cases we tested. Again, the number of 'not sure' answers was fairly low.

3.3. STUDY 3: MEETING THE SUBSTITUTION OBJECTION

It seems that the results of Study 1 and Study 2 can be explained neither by those advocating that a statement is true in virtue of coherence nor by those in favor of the correspondence theory of truth. Instead, the data suggest that both aspects, that is, coherence and correspondence with reality, play a role when people judge a statement to be true or false in the empirical domain. This seems particularly worrisome for correspondentists, as well as scope pluralists arguing for a correspondence view in the empirical domain. After all, it is these theorists who often claim that the correspondence theory captures the ordinary meaning of 'true' in the empirical domain.

Now, correspondentists are likely to make the following objection: the results only provide a challenge to the correspondence theory if the experiments track the folk concept of *truth*. However, so the objection, our experiments failed to investigate the concept of truth and instead prompted

readers to answer a different but related question. In other words, our participants – in their minds – substituted a question about truth with a question about a related concept, for example, truthfulness or truth-for-ness. This phenomenon is known in the psychological literature as ‘substitution heuristic’. Such substitutions often happen, when the substituted question is easier to answer than the original question. Applied to the case in Study 1, evaluating whether Robert’s and Maria’s answers were true is relatively difficult, because Robert and Maria had reliable information that turned out to be misrepresenting reality. Thus, a question that is easier to answer might have been substituted by the participants in our studies. We think that the substitution objection can be cashed out in the following four ways.

According to the first version, when participants stated that Maria’s answer was true, they wanted to convey that she was honest and not telling any lies, or in other words, that she was *truthful*. As the concepts of *truth* and *truthfulness* are quite strongly related, we consider this version of the objection to be particularly powerful. Second, instead of answering a question about the truth of an answer, participants might have wanted to tell us that the protagonist was permitted to give the answer in question, or, more technically, that the protagonist satisfied the norms for making an assertion. The plausibility of this objection depends, of course, on what the norm of assertion is. Some philosophers, for example, Douven (2006), Kneer (2018), Marsili and Wiegmann (2021), and Reuter and Brössel (2018), argue (and present empirical evidence) that the norm of assertion is justified belief. Consequently, if participants follow such a norm and interpreted the question about truth as a question about the norm of assertion, more answers in line with the coherentist prediction are to be expected. Third, participants might have substituted the question about the truth of Maria’s answer by a question about Maria’s epistemic situation, specifically whether she had sufficient epistemic reasons for believing that her answer was true. If that happened, so the objection continues, the results of our study are not relevant for the participant’s concept of *truth*, but rather for what they take to be criteria of truth (see Rescher, 1973, for a discussion of this distinction). According to the fourth and last proposal, when subjects selected the ‘true’ response, they did so, because they believed that the answer was *true for* Maria, not that the answer was true in general.

We agree that all versions of the substitution objection raise serious worries about the correct interpretation of our results. Also, note that every one of the four alternative interpretations of the question would boost the amount of answers that support the coherentist prediction. Given the significance of the objection, we decided to tackle it empirically in three different ways.

First, to avoid the substitution heuristic, it is a common strategy in experimental work to ask the putative replacement question before the original question is asked. Thus, when we first ask participants whether the protagonist of the scenario answered the question to the best of her knowledge, it is

unlikely that the participants will still substitute the subsequent truth question with the same question they just answered. Asking participants whether the protagonist answered the question to the best of her knowledge provides the means to deal with the first three versions of the objection. A person who answers a question to the best of her knowledge (a) is truthful, (b) satisfies the justified belief norm of assertion, and (c) answers in line with the epistemic reasons available to her.

Second, to investigate whether people merely think that the answer was *true for Maria*, we presented participants with the same vignette but then asked participants whether Maria's answer was correct. Whereas the expression 'true for someone' is quite common, a similar phrase 'correct for someone' is not. The Corpus of Contemporary American English (COCA) lists the phrase 'true for you' 171 times; for example, 'True for you, but not for me', but only 10 hits for 'correct for you'. Admittedly, the concept of correctness does not quite match the concept of truth. We therefore need to be careful not to put too much weight on the results of this condition.

Third, to investigate whether any of the four substitution heuristics took place, we did ask the participants to explain their answers. If most participants substitute the original question with a question about, for example, truthfulness or true-for-ness, this should be reflected in people's responses. If, on the other hand, many participants indeed entertain a notion of truth that is guided by coherence, then people should explain their responses accordingly.

3.3.1. *Methods*

A total of 141 participants were recruited on Amazon Mechanical Turk. Four participants had to be excluded for indicating that they were not native English speakers. Of the remaining 137 participants, there were 65 women (one indicating non-binary gender) with an average age of $M_{age} = 34.83$ ($SD = 10.38$). Given the difficulties of ascertaining incoherence that we mention in the discussion of Study 2, we decided to use the same vignettes as in Study 1 and limited ourselves to the Rolex case; that is, participants received the following vignette:

[Rolex] Maria is a watch collector. She keeps all her watches in a safe and knows her collection really well. One day, her friend John asks her, whether she has a 1990 Rolex Submariner in her safe and, if so, could show it to her. Maria answers that she has got a 1990 Rolex Submariner in her safe. After all, she had purchased that watch a few years ago. When Maria opens the safe a little later, she finds out that a burglar has stolen several watches, among them the 1990 Rolex Submariner.

All participants were randomly assigned to three conditions. In the first condition, participants received the Rolex vignette and directly answered the truth question; that is,

[Con1] Was Maria's answer true or false?

Given the low responses to the 'not sure' response, we presented people with only two options: (1) true and (2) false.

This condition also served as a partial replication of Study 1 and as a contrast condition for Conditions 2 and 3. After the question was answered, we did prompt the participants to explain their response to the truth question. In the second condition, participants were first asked a control question; namely,

[Con2] Did Maria answer the question to the best of her knowledge?

People were presented with two options: (1) yes and (2) no.

After the control question, people were first given the truth question and were then asked to explain their response.

In the third condition, people were asked to answer the following correctness question:

[Con3] Was Maria's answer correct?

We presented people with two options: (1) yes and (2) no. Again, people were asked to explain their response.

3.3.2. *Results and discussion*

The results are displayed in Figure 3. In the first condition, in which no control question was asked before the truth question, 71.1% of the participants chose the 'true' response in line with the coherentism hypothesis. In the second condition, a question was placed beforehand as a control to prevent a possible substitution. All participants passed the control question and chose the 'yes' option in response. The percentage of true responses to the truth question in Condition 2 was 51.1%. A Pearson's χ^2 test yielded a marginally significant difference between the two conditions: $\chi^2 = 2.97$, $p = 0.085$. In the correctness condition, 44.7% of the participants selected the 'yes' response. A Pearson's χ^2 test revealed a significant difference between the correctness condition and the original condition: $\chi^2 = 4.12$, $p = 0.032$.

In Condition 1, six out of 31 participants who claimed that Maria's answer was true used words like 'lie', 'deception', or 'truthfulness' in their explanations. It is indeed likely that these participants substituted the truth question with a question about truthfulness. In Condition 2 (with control question beforehand), 22 participants chose the 'true' response. Of these 22 participants, four people gave explanations that indicate that they considered the question to be about truthfulness.

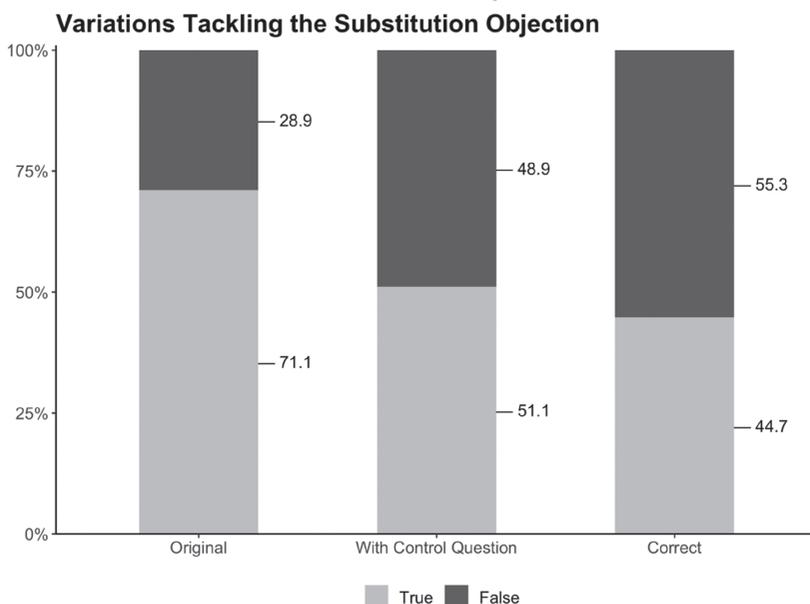


Figure 3. Responses in % for the three conditions in Study 3.

In all three conditions, a substantial number of participants (Condition 1: 14 participants; Condition 2: 8 participants; and Condition 3: 12 participants) provided explanations that indicate the importance of coherence with the available set of beliefs. Some of the explanations that were given by participants were ‘It was true based on the information she had at the time’, ‘Maria was telling the truth based upon the information she knew at the time’, ‘It is true based upon the reports’, ‘Maria answered based on knowledge she knew to be true’, ‘Maria gave an answer based on information they believed to be true at the time’, ‘To the best of their knowledge these answers were true’, ‘She told the truth based on what she knew at the time’, and ‘According to her knowledge, the answer was true’.

When a control question was placed before the test question, a substantial number of participants, 51.1%, considered Maria’s answer to be true. The drop in the true responses are only marginally significant, but a 20% difference indicates that there might well be an effect of the control question on people’s responses to the truth question. A similar case might be made in regard to the correctness results. There was a significant drop of coherentist responses when the ‘true’ question was substituted with a ‘correct’ question. Nonetheless, and importantly, the results of testing the truthfulness and true-for-ness objection have not undermined the main outcome of Studies 1 and 2. Both hypotheses, that is, correspondencism and coherentism, do not find empirical support from our studies.

3.4. GENERAL DISCUSSION OF THE EMPIRICAL DATA

All in all, our studies show four important points about the everyday use of 'true' within the empirical domain.

First, there is substantial variation in the use of 'true'. No uniform answers were given to the question 'Was Robert's/Maria's answer true?' Instead, participants largely disagree on whether Robert's/Maria's answer was true or false. Overall, we tested participants' views in eight different conditions. This variation is a challenge for both monistic and scope pluralistic theories of truth, which both predict little variation in the answers to our cases. In contrast, the variation we found indicates that 'true' is ambiguous since the responses were in fact divided between answers in line with coherentist predictions and answers in line with correspondentist predictions.

Second, vagueness is hardly an issue. The relatively small number of 'not sure' answers suggests that the vignettes did not present borderline cases. Of course, this interpretation can only be applied to questions of vagueness regarding the cases we studied. For a general assessment of the vagueness of *truth*, much more extensive studies would be needed.

Third, conflating the question about truth with a question about truthfulness or about the norms of assertion is a plausible explanation for some of the answers, but it cannot explain the entire variance we found. Instead, we now have even stronger reasons to be confident that the term 'true' is indeed ambiguous. Even when a possibly substituted question was asked before the actual test question, over 50% of the participants gave answers that are best interpreted in accordance with the coherentist picture.

Finally, it is implausible that people are simply muddled in their use of 'true'. The explanations they give for their answers rather suggest that they gave their answers based on a clear understanding of the question.

What can we learn from these findings? The results seem to challenge both monism and some versions of pluralism about truth. Exploring the consequences of these findings for those accounts requires an in-depth discussion we cannot undertake here. In what follows, we will focus on another perspective: In which ways do our findings motivate re-engineering *truth*?

4. *The Project of Re-engineering Truth*

Our experiments exemplify a specific way how empirical research can contribute to projects of conceptual re-engineering: it can be used to test key assumptions of available or newly undertaken re-engineering work.⁹ Not only Tarski's classical explication of *truth* (Tarski, 1933) but also most attempts at giving an exact account of truth rely on the assumption that there is a uniform use of 'true' in everyday language, at least in the empirical domain.¹⁰ Our experiments strongly challenge this assumption and show that research

on coherence may actually be much more relevant to explicating the everyday concept of *truth* than previously thought. Hence, there is reason to re-examine existing theoretical work on truth and especially to reconsider the potential use and range of application of exact accounts of truth. In Section 4.1, we discuss different ways in which *truth* may be re-engineered before we turn to questions of implementation in Section 4.2.

4.1. GOALS AND STRATEGIES FOR RE-ENGINEERING *TRUTH*

The main finding of our experiments is the substantial variance in the everyday use of ‘true’, which indicates that ‘true’ is ambiguous between a correspondentist and a coherentist reading.¹¹ Our results also show that the variance is not an effect of the alleged vagueness of ‘true’. Nor can it be explained away by the conjecture that respondents failed to distinguish *truth* from *truthfulness* and *true for*. All in all, we have good reason to take serious the differences we found in the use of ‘true’.

Our findings also imply that the potential problems sparked by different uses of ‘true’ – ranging from misunderstandings to deliberate deceptions – cannot successfully be dealt with straightforwardly by more fact checking alone, because we also need to ask in which sense ‘true’ is used. Given the fundamental importance of *truth* in science, political discourse, legal discourse, everyday exchanges, philosophy, and so on,¹² the ambiguity of ‘true’ rather makes *truth* a target for conceptual re-engineering. Of course, we need to be careful not to overgeneralize our results from three studies (and two different types of scenarios) to the empirical domain more generally. However, the relatively high robustness of the results and the intuitively high representativeness of the cases make it likely, albeit not guaranteed, that the issue applies widely to everyday and political discourse.

That we advocate taking everyday language serious does not imply that we think the current use of ‘true’ should be left untouched or that a project of conceptual re-engineering should give priority to capturing everyday language as faithfully as possible. Rather, we want to underline that without knowing what is going on in everyday language we do not really know what conceptual re-engineering has to offer and in which ways it can be used to improve upon current everyday language use. Once we realize that ‘true’ is used in both a correspondentist and a coherentist sense, we see that it will not do to confine attention to just one type of use and simply ignore all other uses.

Here are three more promising strategies:

- 1 The *monistic strategy* opts for one target concept (either correspondentist or coherentist) which is supposed to be adopted for some uses of ‘true’, whereas all other uses are not simply ignored but deemed incorrect and actively discouraged.

- 2 The *multidimensional strategy* introduces a single multidimensional target concept with a correspondence and a coherence dimension.
- 3 The *two concepts strategy* advocates for two target concepts, one correspondentist and one coherentist.

The choice of an adequate re-engineering strategy needs to be guided by the goals of the re-engineering project and its intended domain of application. The project we are focusing on aims at addressing *truth* in public discourse and everyday contexts, and it is motivated by concerns that fit with what one may call an ‘agenda of enlightenment’. In particular, we are interested in tools that could help to improve discourse and reasoning. On a most general level, by re-engineering *truth*, we hope to reduce sources of misunderstanding and confusion. More specifically, we may also aspire to foster (self-)critique of false and muddled appeals to, for example, ‘truth’ based on ‘alternative facts’ or an era of ‘post-truth’ or ‘post-fact’.

These purposes have consequences for the re-engineering project. First, target concepts should not primarily be judged by the criterion of similarity to the actual use of ‘true’. Second, the re-engineering project should not take the form of an entirely theory driven explication relying on an overly technical or formal background, but seek target concepts and expressions that could actually find their way into everyday language use.

On this basis, we argue that Strategy (3) is preferable to the alternatives. At first sight, however, introducing more than one target concept for *truth* may seem implausible and far-fetched. Is not Strategy (1) much more attractive? After all, it seems to underlie the most prominent view in the philosophical tradition, which focuses on the use of ‘true’ that can be explicated in the correspondentist fashion. Moreover, Strategy (1) also includes rejecting other uses of ‘true’ as incorrect. This makes it a highly revisionist strategy with respect to the usage of ‘true’ documented in our experiments. As reactions to presentations of our studies have shown, this revisionist element is attractive to many philosophers. It seems that those surprised or shocked by our results tend to think that non-correspondentist uses of ‘true’ in everyday language should just be rejected as incorrect. But although going for a revisionist approach is not problematic in itself, there is little reason for criticizing considerations of coherence *as such*. Denying that coherence is a legitimate and relevant point is just implausible, independent of whether one thinks that coherence should not be associated with the word ‘true’ or that we do not have a satisfactory explication of *coherence*. This, however, means that coherentist and correspondentist uses of ‘true’ are important and should be targets of conceptual re-engineering – especially if one thinks that coherentist uses of ‘true’ should be eliminated.¹³

For those who think that introducing two target concepts will threaten the unity of *truth* in a problematic way, re-engineering *truth* as a

multidimensional concept (Strategy 2) may seem to provide an elegant option. Instead of selectively disfavoring one use of 'true' (as in Strategy 1), we could insist on there being just one concept of truth but nonetheless accept correspondence and coherence as genuine aspects of truth. This would be a notable advantage if we thought that we should be able to handle situations in which correspondence and coherence diverge by an overall judgement of truth which combines consideration of both aspects. However, a multidimensional concept would confront us with the question of what to say about the considerations that determine the weighing needed for arriving at an overall judgement. It is far from clear that any fairly determinate answer can be made plausible. But most importantly, the very possibility of amalgamating considerations of correspondence and coherence in an overall judgement of truth seems not desirable at all, given our goal of reducing sources of misunderstanding and confusion.

Strategy (3) avoids these problems and provides an appropriate and elegant way of addressing the finding that in everyday discourse correspondence with facts and coherence are both important and both covered by the concept of *truth* as reflected in the use of the word 'true'. Since both correspondence and coherence are perfectly legitimate concerns, an engineering project had better not privilege the correspondentist at the expense of the coherentist use of 'true', or the other way around. Clarity should rather be gained by explicitly distinguishing correspondentist and coherentist uses. This speaks in favor of introducing specifically correspondentist and coherentist concepts, in line with Tarski's (1944, pp. 355–356) advice. The good news is, of course, that philosophers have already done a great deal of work by distinguishing correspondence from coherence and other related concepts such as *justification* and *confirmation*. Applying these distinctions to everyday discourse and reasoning should therefore be helpful to draw attention to different criteria of truth: empirical evidence which verifies or falsifies a statement by showing that it does (not) correspond with the facts in the case of correspondence and support by basic commitments and background assumptions in the case of coherence.

4.2. IMPLEMENTING TWO TARGET CONCEPTS

How could such a re-engineering project be implemented? What is probably most important to note first is that the philosopher's concepts of correspondence and coherence are *both* target concepts for *truth* as it is manifest in the use of 'true' in everyday language. If our analysis is correct, it would be a mistake to think, as proponents of the monistic strategy typically do, that *correspondence with facts/reality* (or some more technical explication thereof) is the single most appropriate target concept that could serve as a more clearly delineated replacement for *truth* in everyday discourse.

So the first point a conceptual re-engineer of the everyday concept of *truth* should insist on is just that we need to promote the distinction between correspondence and coherence. If ‘true’ is used in an everyday context, the resulting utterance may well be in need of disambiguation. It is therefore always a good idea to ask yourself or the speaker whether it is correspondence with facts she has in mind or rather coherence with some set of statements – especially so if there is a suspicion that the two considerations may actually pull in different directions.

The second point concerns terminology. The analysis we defend also speaks against using ‘true’ for one of the target concepts. This point might be met with skepticism. Is discouraging or even eliminating the use of ‘true’ in everyday language not quite an unrealistic proposal? And, surely, introducing technical terms like *true_{coh}* into everyday language is a silly idea, isn’t it? Nonetheless, we think that it is often worthwhile to avoid the word ‘true’ and its cognates and look instead for expressions that are available as means of clarification. So we should, for example, ask questions such as ‘*Is it a fact that France is in turmoil because of the Paris accord? Or is it a claim that fits with Trump’s beliefs?*’ And confronted with the Rolex case, it would be helpful to say something like ‘Maria’s answer *was coherent* with what she knew, but *the fact is* that the 1990 Rolex Submariner was not in her safe.’

Both points might seem rather dull insofar as the general advice to avoid ambiguous expression where they could lead to confusion seems rather uncontroversial. But the application of this advice to the use of ‘true’ in everyday discourse is not uncontroversial at all. It is neither available to philosophers who advocate the monistic strategy nor to people who think that more fact checking alone will cure those who call claims not in line with facts ‘true’.

Moreover, the potential for clarification does not stop at pointing out ambiguous uses of ‘true’. Careful conceptual work has more to offer than the mere distinction between the concepts *correspondence* and *coherence*. Best known, of course, is the extensive work which has been done on developing concepts of *truth* that avoid Liar and Revenge paradoxes. But for the project we are promoting here, these problems are not crucial. Rather, some very elementary issues a conceptual re-engineer needs to address offer considerable potential for clarification of everyday discourse¹⁴: For example: How many argument places should the target concept have? What should be its range of meaningful application? Should it be a classificatory or a comparative concept? Used in the sense of correspondence, ‘true’ may be constructed as a classificatory one-place predicate intended to be applied to sentences, beliefs, statements, utterances, and the like. The concept of *coherence*, however, has a different structure (which probably varies from context to context, but we ignore this complication here). We need at least a two-place predicate that expresses a relation between an individual sentence (or belief or . . .) and a specific set of sentences (or a set of beliefs or . . .), and this relation admits

of degrees. So whereas in the correspondentist reading ‘Maria’s answer is true’ can be re-engineered as ‘Maria’s answer *corresponds with the facts*’, in the coherentist reading, it is elliptical for something like ‘Maria’s answer *coheres perfectly with the beliefs she held at the time of giving the answer.*’ All this, we submit, has a critical potential, which has not been exploited enough for clarifying uses of ‘true’ so far.

Finally, we need to keep in mind that it makes little sense to design projects of conceptual engineering as addressing one specific concept only. As soon as we start re-engineering *truth* in terms of correspondence or coherence, further concepts such as *fact*, *belief*, and *correspond* come into focus and potentially become targets for conceptual re-engineering. All the points just mentioned are familiar from the many projects that have aimed at explicating *truth* or *coherence*. We emphasize them here, because they all have great potential to contribute to the clarification of everyday use of ‘true’ – a fact that is lost on those who simply assume that in empirical matters the everyday concept of *truth* just boils down to correspondence.

5. Conclusion

Empirical research can be used to test and challenge (or even refute) relevant assumptions about the actual use of a concept. In everyday language, the use of ‘true’ shows substantial variance within the empirical domain indicating that ‘true’ is ambiguous: in several scenarios, some subjects used ‘true’ in line with correspondentist, others in line with coherentist predictions. Given the results of our studies, we suggested to re-engineer the concept of truth. More specifically, we argued that in order to improve public discourse and reasoning, we need at least two target concepts, one capturing truth in the correspondentist sense and one covering coherence. Because truth theorists so far have focused on debating explications at the expense of careful and empirically informed studies of the concepts actually used, they have considerably underestimated how much systematic critical potential of clarification their work has.

6. Acknowledgments

Previous versions of this paper have been presented in Bern, Bochum, Prague, St. Andrews, and Zurich. We would like to thank the audiences and in particular Gregor Betz, Mark Pinder, Joseph Ulatowski, Pascale Willemsen, and anonymous referees for this journal for their very helpful feedback. Open Access Funding provided by Universitat Zurich.

7. *Funding*

This research was supported by the Swiss National Science Foundation (SNSF) (PCEFP1 181082).

Kevin Reuter
Department of Philosophy
University of Zurich

Georg Brun
Institute of Philosophy
University of Bern

NOTES

¹ 'Post-truth' was selected as Word of the Year 2016 by Oxford dictionaries; see <https://languages.oup.com/word-of-the-year/word-of-the-year-2016> (accessed 20-08-2020).

² One might think that it is more plausible that 'true' exhibits some form of polysemy rather than ambiguity (see Lanius, 2019, for this distinction). Since the difference between ambiguity and polysemy plays no role for our purpose, we simply speak of ambiguity throughout.

³ A note on terminology: we use 'conceptual engineering' to refer to the explicit and intentional development of a concept in order to solve some theoretical or practical problem and 'conceptual re-engineering' for conceptual engineering that is guided by a concept in use and simultaneously aims at replacing this concept.

⁴ For a thoroughgoing defense of truth pluralism based on several empirical studies on the folk concept of truth, see Ulatowski (2017).

⁵ In that sense, our philosophical project is also an invitation for those skeptical of our hypotheses to suggest ways of operationalizing their notion of truth for empirical studies.

⁶ We did not opt for measuring people's responses with a Likert scale, because graded response options would arguably bias people into thinking more in a coherentist way. Correspondentists usually consider truth to be binary; that is, Maria's and Robert's answers are either true or false, but not something in between.

⁷ We would like to thank a reviewer for this journal for pressing us on this point.

⁸ Scenarios that pitch correspondence against incoherence are more difficult to design than vice versa. Especially the reliance on memory is problematic because it is not clear whether a statement that is incoherent with a forgotten piece of information yields the right kind of incoherence that is needed. Thus, some people might have selected 'true' despite holding a coherentist conception of truth, because they believed Robert's or Maria's answer to be coherent. Another way to realize the incoherence of an answer with a certain set of beliefs would be to use cases of dishonesty. However, such cases might bias people into interpreting the question to be about truthfulness instead of truth (see also Section 3.3).

⁹ For discussions of other contributions of empirical research to conceptual re-engineering, see Shepherd and Justus (2015), Machery (2017, ch. 7), Pinder (2017), and Koch (2019).

¹⁰ 'Explication' in Carnap's (1962) sense is a method of conceptual re-engineering which seeks exact and fruitful concepts for some theoretical purpose (see Brun, 2016). Hence, most

theories of truth which aim at giving an exact account of truth can be seen as proposing explications of *truth* and therefore as projects of conceptual re-engineering.

¹¹ As in our empirical studies, we simplify by focusing on correspondentist and coherentist readings of ‘true’ although it is quite possible that additional uses may be important in everyday language.

¹² This is not to say that science aims at relying on true statements only (think of idealizations), similarly for other discourses.

¹³ Note that our objection specifically targets the application of Strategy (1) to the project we sketched in the preceding paragraphs. In another setting, focusing exclusively on correspondentist or on coherentist uses of ‘true’ may be an adequate strategy.

¹⁴ Brun (2016) contains a more comprehensive list of questions explicators have to address.

REFERENCES

- Barnard, R. and Ulatowski, J. (2013). ‘Truth, Correspondence, and Gender,’ *Review of Philosophy and Psychology* 4(4), pp. 621–638.
- Barnard, R. and Ulatowski, J. (2016). ‘Tarski’s 1944 Polemical Remarks and Naess’ ‘Experimental Philosophy,’ *Erkenntnis* 81(3), pp. 457–477.
- Bourget, D. and Chalmers, D. (2014). ‘What Do Philosophers Believe?’ *Philosophical Studies* 170, pp. 465–500.
- Brun, G. (2016). ‘Explication as a Method of Conceptual Re-engineering,’ *Erkenntnis* 81(6), pp. 1211–1241.
- Carnap, R. (1962). *Logical Foundations of Probability*, 2nd edn. Chicago/London: University of Chicago Press/Routledge and Kegan Paul.
- Douven, I. (2006). ‘Assertion, Knowledge, and Rational Credibility,’ *Philosophical Review* 115, pp. 449–485.
- Edwards, D. (2012). ‘Pluralist Theories of Truth,’ in *The Internet Encyclopedia of Philosophy*. University of Tennessee at Martin.
- Kneer, M. (2018). ‘The Norm of Assertion: Empirical Data,’ *Cognition* 177, pp. 165–171.
- Koch, S. (2019). ‘Carnapian Explications, Experimental Philosophy, and Fruitful Concepts,’ *Inquiry* 62, pp. 700–717.
- Kölbl, M. (2008). ‘‘True’ as Ambiguous,’ *Philosophy and Phenomenological Research* 77(2), pp. 359–384.
- Lanius, D. (2019). *Strategic Indeterminacy in the Law*. New York: Oxford University Press.
- Lynch, M. (2012). ‘Three Questions About Truth,’ in N. J. Pedersen and C. D. Wright (eds) *Truth and Pluralism: Current Debates*. New York: Oxford University Press, pp. 21–41.
- Machery, E. (2017). *Philosophy Within Its Proper Bounds*. Oxford: Oxford University Press.
- Marsili, N. and Wiegmann, A. (2021). ‘Should I Say that? An Experimental Investigation of the Norm of Assertion,’ *Cognition* 212.
- Næss, A. [original: Ness](1938). *‘Truth’ as Conceived by Those Who Are Not Professional Philosophers*. Oslo: Jacob Dybwad.
- Næss, A. (1953a). *An Empirical Study of the Expressions ‘True’, ‘Perfectly Certain’ and ‘Extremely Probable’*. Oslo: Jacob Dybwad.
- Næss, A. (1953b). *Interpretation and Preciseness. A Contribution to the Theory of Communication*. Oslo: Jacob Dybwad.
- Pinder, M. (2017). ‘Does Experimental Philosophy Have a Role to Play in Carnapian Explication?’ *Ratio* 30(4), pp. 443–461.
- Putnam, H. (1981). *Reason, Truth and History*. Cambridge: Cambridge University Press.
- Rescher, N. (1973). *The Coherence Theory of Truth*. Oxford: Clarendon Press.

- Reuter, K. and Brössel, P. (2018). 'No Knowledge Required,' *Episteme* 16(3), pp. 303–321.
- Scharp, K. (2013). *Replacing Truth*. Oxford: Oxford University Press.
- Shepherd, J. and Justus, J. (2015). 'X-Phi and Carnapian Explication,' *Erkenntnis* 80(2), pp. 381–402.
- Tarski, A. (1983)(1933). 'The Concept of Truth in Formalized Languages,' in *Logic, Semantics, Metamathematics*, 2nd edn. Indianapolis: Hackett, pp. 152–278.
- Tarski, A. (1944). 'The Semantic Conception of Truth and the Foundations of Semantics,' *Philosophy and Phenomenological Research* 4, pp. 341–375.
- Turri, J. (2013). 'The Test of Truth: An Experimental Investigation of the Norm of Assertion,' *Cognition* 129, pp. 279–291.
- Ulatowski, J. (2017). *Commonsense Pluralism About Truth. An Empirical Defence*. London: Palgrave Macmillan.
- Wright, C. (2005). 'On the Functionalization of Pluralist Approaches to Truth,' *Synthese* 145(1), pp. 1–128.
- Young, J. (2001). 'A Defence of the Coherence Theory of Truth,' *The Journal of Philosophical Research* 26, pp. 89–101.