

Modeling Truth

Paul Teller
prteller@ucdavis.edu

Draft March 2016

| Key words: Truth, accuracy, precision, models, pragmatism, common ground.

Modeling Truth

Abstract

Many in philosophy understand truth in terms of precise semantic values, true propositions. Following Braun and Sider, I say that in this sense almost nothing we say is, literally, true. I take the stand that this account of truth nonetheless constitutes a vitally useful idealization in understanding many features of the structure of language. The Fregean problem discussed by Braun and Sider concerns issues about application of language to the world. In understanding these issues I propose an alternative modeling tool summarized in the idea that inaccuracy of statements can be accommodated by their imprecision. This yields a pragmatist account of truth, but one not subject to the usual counterexamples. The account can also be viewed as an elaborated error theory. The paper addresses some prima facie objections and concludes with implications for how we address certain problems in philosophy.

1. Methodological preamble. One often starts from the presumption that we should look for the one right, the TRUE account of truth. The methodology of the present paper supposes that, as in dealing with complex subjects in science, we can learn about a subject by fashioning idealized¹ accounts with the expectation that complementary idealized accounts will be forthcoming, any of which are of value insofar as they provide illuminating understanding of a robust body of phenomena.²

One such account, what I will refer to as “referential semantics”, works in terms of propositions - precise, determinate things formed in terms of more

¹ I use ‘idealize’ and ‘idealization’ very broadly to cover any inaccuracy where the implication is that the inaccuracy won’t spoil the intended application.

² See my (2001, 401) for the illustration of complementary hydrodynamic and statistical mechanical accounts of water.

specific determinate semantic values such as objects, properties, and relations. For example, the proposition that John's height is 180cm is true just in case the referent of 'John' has the property expressed by 'height is 180cm'; and where a statement is true just in case it expresses a true proposition. In this simple example, the truth condition for the statement that John's height is 180cm is that the referent of 'John' has the property attached to 'height is 180cm'.

Referential semantics starts by characterizing lexical entries as attached to things in the world, referents, properties, and relations.³ I invite readers to consider the possibility that this characterization is a simplification, an idealization. Idealization is the workhorse of science, and can likewise be in philosophical analysis, but, as in science, we must stay vigilant about the limits of an idealization's successful application. The position I take in this paper is that the idealizations of referential semantics work splendidly when it comes to recursive and other features of language, but that it proves problematic when it comes to understanding how language gets connected to the world. Since referential semantics' account of truth rests on its account of the word-world connection, the ways in which that connection is idealized devolve upon referential semantics' account of truth.

A point to emphasize: The fact that I urge an alternative account of truth (highly idealized in its own ways!) does not mean that I urge that we stop using that of referential semantics. For almost all practical applications, the informal counterpart of referential semantics' truth proves the best, as I will explain in some detail in section 11. But for a finer grained understanding of the success of human language in its representational function we may gain with a complementary model. In addition, as I will mention briefly at the end, the more finely grained approach to truth suggests fruitful ways of approaching some other important issues in philosophy.

Another way to express the idealization intrinsic to referential semantics is that it idealizes away from the ubiquitous phenomenon of vagueness, in which terms we reexpress the worries of the last paragraph with the question, what is it for a vague statement to be true? Indeed, deidealizing truth and asking what it is for a vague statement to be true may be seen as two sides of the same problem. In this paper I will approach the problem entirely from the point of view of deidealizing truth. Given the suggested duality between deidealizing truth and truth for vague statements, the present account begins an account of vagueness, but only begins as it will not discuss ancillary puzzles such as higher order vagueness or sorites arguments. In a companion paper (to appear b), I will examine the close connections between the present results and their application

³ Henceforth "properties" will be understood as "properties and relations".

to some other issues about vagueness.

2. Terminology. Another way to state the idealization of referential semantics is that it represents statements as completely precise – the semantic values are thought of as attached to specific referents and properties; and then to be true a statement is taken to be completely accurate – in the simplest cases, to be true is for the referent to have the property. Conversely, the idealized limitations of referential semantics have to do with imprecision and inaccuracy. A word about how I will understand these terms.

My use of ‘imprecision’ differs from the way it is often used in the sciences. Suppose I ask, what integer am I thinking of: Mary guesses that it’s in the interval [3, 43], Bob guesses that it’s in [5, 17], and Martha guesses that the number is 12. In common scientific usage Mary’s guess counts as less precise than Bob’s, and Martha’s guess is completely precise. In my usage all three statements of the guesses count as completely precise because all three pick out a completely determinate proposition. In my intended usage of ‘imprecision’ an expression counts as imprecise to the extent that it fails in specifying a completely determinate intended referent, property, or proposition; and this even when all contextual factors have been taken into account, eliminating ordinary ambiguity, filling in values for implicit parameters, referents assigned to indexicals, and the like. Thus in my usage the statement, *it’s the number you would get if you counted for a minute*, would count as imprecise.

What is the relation between the notions of imprecision and vagueness? Vagueness is generally introduced in one of two ways: i) Facilitating construction of a sorites paradox (“being sorites susceptible”); ii) Having borderline cases. Do these coincide? I need not settle the issue. I take ii) to be at least as wide as i) and I use ‘imprecise’ for being vague in this (possibly) broader sense. ‘Imprecise’ also has the advantage of supporting the antonym, ‘precise’⁴

Clearly precision comes in degrees; one statement can be more or less precise than another. When we say that a statement is precise, full stop, we often mean that it is precise enough so that further refinements in precision would be irrelevant to present purposes.

I am going to claim that, with possible exceptions in finite mathematics, no statement is both completely precise and completely accurate. How is (in)accuracy to be understood here? Within the framework of referential semantics to say of a statement that it is completely accurate is just to say that it is true. Then to say of it that it is inaccurate is to say that it is false. But in addition “inaccurate” carries at least the implication that the way in which the

⁴ In discussions of vagueness, many use (im)precise as I do. For example, Williamson (1994) uses “precise” as the antonym of “vague” throughout his book.

statement is false admits of correction by adjustment. Thus, it would be misleading to apply the evaluation “inaccurate” to the statement that it is day made in the middle of the night but appropriate to characterize as inaccurate the statement that it is twilight.

Since I will be proposing an alternative to idealized referential semantics I need to get the spirit of the last example in some other way. When I say of an imprecise statement that it is also inaccurate this will be to indicate that the statement can be corrected, can be brought into closer conformity with the way things are, without also making the statement more precise. Suppose that using the sewing drawer tape we have measured John’s height at 180cm. Then I increase the precision of the statement that John is tall by saying that John’s height is close to 180cm. But if I make the inaccurate statement that John’s height is close to 170cm I increase the accuracy but not the precision, with the correction that his height is close to 180cm.

One should always keep in mind that a statement can be accurate in some ways and less so in others. Conjunctions provide the easiest but far from the only examples. Also, in the same way as precision, accuracy comes in degrees. One statement can, in a given respect, be more accurate than a second while still admitting of further correction. When we say that a statement is accurate, full stop, we usually mean that the statement is accurate enough so that further refinement in accuracy isn’t relevant to contextually specified interests.

Finally, I will reserve ‘inexact’ as a cover term for ‘imprecise and/or inaccurate’, and likewise ‘exact’ for ‘both precise and accurate’.

I will talk about statements, where by ‘statement’ I will intend our preanalytic idea of “what is said”, but without presupposing that what is said is something completely precise. To differentiate between sentences (expressions, utterances) and statements I will use single quotes to talk about sentences, expressions, and utterances, and italics to talk about statements; and I will treat sentential fragments similarly. I will reserve double quotes for any of the many uses of quotation other than explicit mention of terms.

3. The poverty of exact truths in science. Most of us suppose that science provides us with exact truths. A plausible line of reasoning supports this impression: We think that we can discover determinate natural kinds and quantities, that we can directly attach these to words, in terms of which we can then formulate perfectly exact singular statements and general laws. I am skeptical.

I will illustrate with the special but particularly plausible case of fundamental physics. All fundamental theories in physics are idealizations⁵, in particular idealizing the nature of the objects and quantities that these theories study. Take, for example, the quantity, mass. Newtonians thought that the term 'mass' was univocal, referring to a perfectly determinate, unique natural quantity. But relativity blurred the distinction between mass and energy. In the general theory of relativity, gravitational mass/energy isn't even exactly localizable; and relativity, both special and general, are idealizations. Quantum theories further complicate how to think about mass as mass serves as a renormalization parameter: Mass is a relative quantity, relative to the strength of an interaction. These considerations strain the claim that 'mass' has been univocally attached to some specific quantity.⁶

NB: I am not claiming that science provides no truths. For example, I take it to be true that water is H₂O. But the statement that water is H₂O is not a precise statement. What quantity of H₂O will count as a quantity of water? Do we mean 100% pure bulk H₂O, of which there is none, and which, if there were any, would have significantly different properties than the stuff we call 'water'? Do ice and water vapor count as water? At just what temperatures and pressures is H₂O supposed to be water? I am claiming that science provides us with no statements that are both perfectly precise AND perfectly accurate, thus no truths as truth is characterized in referential semantics.

We take science to provide the best of human representation and knowledge. If we have no perfectly precise and accurate statements in science, where else should we expect to find them? Yet much that we say, in and out of science, is true. My objective is to make out a notion of truth on which the statement that water is H₂O, as well as a great many things we say every day, count as true.

4. A more general reason for suspecting that all of our representations fall short of being both completely precise and completely accurate. The world is extremely complicated. To succeed in getting around, our native concepts as well as the ones we self-consciously fashion make all sorts of simplifications not exactly faithful to the facts. As is well known in the case of color,⁷ the concepts that arise through our perceptual capacities can be

⁵ This is supported in my (2004a, 432-3, 436-7)

⁶ For much more on the idealized status of quantity terms in physics see my (to appear a).

⁷ In chapter II of his (2006) Giere provides a compelling summary of the grounds for the conclusion that color perception involves all kinds of shortcuts short of exact reflectancies of external objects.

expected to use simplifications, shortcuts, what are strictly speaking distortions but that work well enough for our needs and more efficiently than would ones that apply with greater accuracy. It is reasonable to expect that distortions that more than compensate in efficiency for what is lost in accuracy will occur throughout our information processing capacities.

That there is ubiquitous distortion in our non-verbal representations suggests that the same sort of thing may go every bit as much for human verbal representation, that verbal representation ubiquitously involves some level of imprecision and/or inaccuracy. I recognize one plausible exception: representation of combinatorial facts, logic, and the like, that we can collectively characterize as finite mathematics. But for mathematics more broadly, as soon as incompleteness results apply, unintended interpretations compromise complete precision.⁸ (Henceforth I will take the exception of finite mathematics for granted and not mention it explicitly.)

5. A very different reason for concluding that virtually nothing we say is, literally, true. In “Vague so Untrue” Braun and Sider (2007) argue that sentences, expressions, or utterances are rarely, if ever, either true or false.⁹ They start from the assumption that

...the properties, relations, and propositions that are candidates for being the meanings of linguistic expressions are precise. (134)

They help themselves to this assumption because it functions as a framework principle, presupposed throughout the enterprise of referential semantics. They continue with a factual claim for which they do not argue, but which many will grant:¹⁰

[T]he facts that determine meaning (for instance, facts about use, naturalness of properties, and causal relations between speakers and properties) do not determine a unique property to be the meaning of ‘red’ [and likewise for expressions very broadly]. (134)

From the two foregoing assumptions they conclude that

⁸ Note that this exception is systematic and fits into the larger scheme developed below: It can be seen as the limiting case in which truth-apt imprecise statements coincide with their idealized semantic alter-ego counterparts. See below.

⁹ As Braun and Sider recognize, this thinking goes back to Frege. See their note 1, p. 133 for some references

¹⁰ And for which I will argue in section 8.

...therefore [there is] no unique proposition that a sentence containing 'red' expresses. (134)

Next they appeal to another unquestioned principle guiding referential semantics:

To be either true or false, a sentence must have a unique meaning [viz., express a unique proposition]. (134)

and conclude:

"So, on our view, utterances are rarely (if ever) true or false. (134)

QED.^{11 12}

6. Questioning Braun and Sider's framework. To summarize their argument:

- 1) Assume that semantic values, in particular propositions, are precise.
- 2) Assume that for a statement¹³ to be true is for the statement to express a true proposition.
- 3) Further factual claim: Statements are always (or almost always) imprecise, in particular having no semantic value.

¹¹ Braun and Sider understand the plainly unproblematic activity of ascribing truth to utterances by saying that "ordinarily speakers typically and harmlessly ignore vagueness." (135), but provide no account of how this works. My positive account of truth will, in section 11, support a developed analog of what they suggest.

¹² I believe that Braun and Sider's problem cannot be resolved by any of the existing accounts of vagueness. To some extent Braun and Sider support this claim in their paper. I will defend this claim in (to appear b) and here continue with the positive approach of showing how their issue can be addressed directly by refashioning how we think about truth.

¹³ I here shift from Braun and Sider's use of 'sentence' and 'utterance' to 'statement' to bring their terminology into line with mine. Since what counts here are the semantic values, nothing turns on whether we take these to be values of sentences, utterances, or statements.

4) From 1) and 3): Statements do not (or rarely) express any proposition.

5) From 2), and 4): Statements are never (or rarely) true (or false).

We should consider turning this use of modus ponens into a modes tolens; Those who accept 1) and 3) but choke on conclusion 5) will reject 2): If one accepts 1) and 3), and then also takes many statements to be true, they must be true in some way other than by expressing a true proposition. Braun and Sider are working within a widely accepted framework according to which 1) propositions are precise and 2) for a statement to be true is for it to express a, true proposition. Thus those who reject conclusion 5) but accept 3) are committed to challenging this framework. Taking 'semantic value' as a term of art for which 1) is presupposed, such a challenge amounts to rejecting 2)

My project is to begin sketching an alternative framework that will show what is involved in truth as we commonly use that notion.

I must emphasize: *I do not advocate giving up the familiar framework that includes presuppositions 1) and 2).* The mystique of Kuhnian paradigm shifts to the contrary notwithstanding, science does not generally discard useful frameworks. Newtonian physics has hardly been abandoned. Rather, it thrives as an extraordinarily powerful theoretical tool, and not just for making predictions: Imagine trying to explain the terrestrial tides with the general theory of relativity! Nor has Newtonian mechanics been degraded to a merely useful framework because it has been superseded by true theories. Relativity and quantum theories are every bit as much idealizations as is Newtonian mechanics. Indeed, in science a wide range of idealized theories function in their own right in helping us to understand the world.¹⁴ Equally, the framework of referential semantics will continue to provide the best way of understanding many important features of language. In particular, my complementary approach will make heavy use of referential semantics' idealized way of understanding truth.

7. Literal truth that accommodates both imprecision and inaccuracy.

How might we think of the truth of a statement other than the expression of a true proposition? Consider the evaluation of maps and their trade-off between imprecision and inaccuracy. Imagine a simple precise map with sharply drawn lines of slightly different length to represent the distances from A to B and from A to C as slightly different. This map is slightly inaccurate: the distances represented are the same.

Now imagine a second map very like the first, except that the representations of distances have been made imprecise by making them fuzzy,

¹⁴ As I argue in my (2004a).

slightly out of focus. This map does not give the distances A-B and A-C exactly. It shows them as close, but it would be hard to determine whether this second map represents the distances as exactly the same or perhaps slightly different. Metaphorically speaking, the inaccuracy of the first map has been “swallowed up” by the imprecision of the second map.

As for maps, similarly for statements. The idea of a completely precise and true statement corresponds to a map with complete precision and accuracy. In practice there are no such maps.¹⁵ And, in practice, no such statements. But where there is no inaccuracy within a statement’s level of precision we will take that absence of inaccuracy to constitute the truth of the statement. In other words, an imprecise statement counts as true just in case increasing the statement’s agreement with the language independent world requires making the statement more precise.¹⁶

To implement this strategy we will retain, as an idealization, statements thought of as precise and thought of as expressing precise propositions. Idealizations are generally false. But they can do the work of exact truths. We will then take a second kind of statement to be imprecise and true by virtue of the work of exact truth accomplished by the corresponding precise but not completely accurate statement. Metaphorically speaking, a statement can be, literal, true by virtue of its imprecision “smoothing over” the inevitable inaccuracy of its precise, idealized counterpart. The strategy will be to see precise but strictly speaking inaccurate statements and imprecise but literally true statements as two sides of the same semantic coin.

8. Semantic alter-egos. Exposition of the analysis will take place in two steps. In the first we simplify by supposing that there is a referent of ‘John’, that there is a physical quantity, height of people, so that in our idealized example the statement, *John’s height is 180cm*, can be understood as picking out the sort of proposition characterized in referential semantics. I simplify here to set ideas. In the next section these simplifications will be let go.

We again consider the case in which we have checked John’s height with the sewing drawer tape and come up with the conclusion: 180cm. There are two ways of understanding the statement, *John’s height is 180cm*. First in the idealized understanding of referential semantics just aired, that there is a person, John, there is a physical quantity admitting of precise values, height in

¹⁵ Compare Russell (1923, p. 68): “All vagueness in language and thought is essentially analogous to this vagueness which may exist in a photograph.”

¹⁶ This initial formulation doesn’t work for tautologies and contradictions, more broadly for truths and falsities of finite mathematics, which in any case fall outside the whole analysis.

centimeters, and that John's height is 180cm. I will summarize this understanding by saying that John's height is EXACTLY 180cm.

But, even waving the worries about the other idealizations, no one's height is EXACTLY 180cm. Nonetheless, where the discrepancy between an idealized exactly 180cm and John's actual height characteristics doesn't matter for current concerns this false¹⁷ statement *functions* as an exact (that is completely precise and completely accurate) truth. It will be important for what follows that functioning as an exact truth be understood as functioning as an exact truth with respect to the needs and interests of concern and that functioning be understood as functioning generally, for any extension of the situation that might reasonably be supposed to come up.¹⁸ When these circumstances obtain we will say that the *conditions of application* obtain for the false statement, *John's height is exactly 180cm*.

But the statement, *John's height is 180cm*, can also be understood as: *John's height is CLOSE AS MAKES NO DIFFERENCE TO 180cm*¹⁹. In making the statement with this understanding and with the other idealizations in play we are supposing that there is a person, John, that John has some completely precise height in centimeters, but we refuse to say exactly what that number is, instead saying that it is some number close enough to 180cm so that there is no practical difference between that number and 180cm.

In the circumstances we are considering this statement is, literally, true. What circumstances? Exactly the foregoing conditions of application for the false *John's height is exactly 180cm*. We want to use the conditions of application for false, precise statements to characterize the conditions under which the imprecise analog will qualify as true.²⁰ For *John's height is close as makes no*

¹⁷ If I say that no completely precise statement is true, won't I have to say, with Braun and Sider, that none are false either? Yes. Use of 'false' here is a special case of the expository simplifications assumed in this section, addressed with the considerations of section 9, and more generally, section 11. Section 9 will direct us to take the statement as *presupposing* that there is some precise number that is John's height in centimeters and inches but that this number is not 180cm. On this reading the statement fails of presupposition and so comes out as truth valueless. But what I say about functioning as an exact truth is not affected.

¹⁸ I will expand on this requirement below.

¹⁹ Compare: The top of my table is flat. Not geometrically flat, but close enough that differences won't matter for present concerns.

²⁰ Below I will explain why conditions of applications are not traditional truth conditions.

difference to 180cm, to be true we must take into account: close enough for what? Clearly, one intends contextually indicated needs and interests. This interpretation corresponds to the way in which we characterized the idea of functioning as an exact truth in our discussion of the false, precise analog.

John's height is exactly 180cm and *John's height is close as makes no difference to 180cm* together provide an example of what I will call *semantic alter-egos*. The implication of 'semantic alter-egos' is that the two statements get the same representational work done, the one as a false precise statement that nonetheless functions as an exact truth, the second as an imprecise statement that is, literally, true, true when the conditions of application for its precise semantic alter-ego obtain, the conditions under which the precise statement functions as a exact truth. By appealing to the idea of semantic alter-egos the "loose' talk" under which one might apply 'true' to the precise statement has been absorbed into the imprecision of the imprecise statement. In this way we have consistently accommodated inaccuracy and truth by shifting the burden of inaccuracy to imprecision. On this account, in the functioning of language inaccuracy and imprecision are intimately connected, adding substance to the old but nonspecific observation that vagueness, rather than a defect, is essential to the operation of language.

How should we understand the statement, *John's height is NOT 180cm tall*? As no ones height is exactly 180cm, with negation the precise semantic alter-ego comes out as true²¹ for all real people. But in the circumstances we are considering the negated precise alter-ego does NOT function as a truth. In ordinary circumstances, to say of someone who is just a millimeter off of exactly 180cm that they are not 180cm would be extremely misleading. So what we need, for statements as they are ordinarily understood, is to ask whether the precise negated alter-ego functions as a truth for us as we ordinarily use language, whether or not the precise alter-ego is, strictly speaking, true or false. This works out exactly right. After the result of 180cm with the sewing drawer tape the precise semantic alter-ego, *John's height is not exactly 180cm*, does not function as a truth, so that the conditions of application fail, and the statement understood as *John's height is not close as makes no difference to 180cm* comes out false.

Objection! Aren't these "conditions of application" just a trivial renaming of "truth conditions" that are, again, precise without qualification? The objector claims that, though perhaps very complicated, there are completely determinate facts about the context of use and the human needs and interests in that context. These facts will then support exact conditions under which a given statement functions as a truth with respect to the needs and interests in the context. Exactly these conditions, then, will be the claimed truth conditions for the

²¹ The considerations of note 17 apply likewise here.

statement in the context in which the statement is made.

No, there are no such completely determinate conditions, if only because human needs and interests are themselves not completely determinate. Consider: Exactly what stuff will satisfy your desire for chocolate ice cream? There will be cases in which you will say “yes”, cases in which you will say “no”, and, most likely, cases in which you will say “I’m not sure”. Should we be automatons, there could be facts of our physical functioning that would determine under exactly which physical conditions we would say which one of these three things. But in some of the cases it will be indeterminate whether or not the response really reflects experienced satisfaction of the desire for chocolate ice cream. From the point of view of our interests, just where the line is drawn will be arbitrary, determined by things that, from the point of view of human interests, are irrelevant accidents. From the point of view of human interests just where the line is drawn in hard cases will be determined case by case by decisions, decisions which in really close cases will be arbitrary, will be “don’t cares”, will be cases that just as well could have gone the other way.

Similar comments apply very generally to the operation of the mechanisms by which we apply language to the world. On the attitude defended here, talk of truth conditions are part of a highly idealized model of the word-world link. The truth conditions for the statement, *the fire engine is red*, presumes a completely specific property, being red, with the truth conditions then being that the fire engine has this property. But application of an idealized model to the world always has wiggle room, wiggle room that is incorporated in the idea of conditions of application’s “functioning as a truth”. The idealized truth conditions are an idealized presentation of standards of correct application of words to the world. When we deidealize²² truth conditions, standards for correct application must be realized in some humanly applicable form, and if I am right about the gap between the complexity of the world and what is humanly realizable, the realization won’t be something that will operate with complete precision. This is just what is captured by the (admittedly schematic) “functions as a truth”.²³

Another good question: Given an imprecise statement, what is its precise semantic alter-ego? The most important point at this level of analysis is that, given a statement understood as imprecise, it need not have a unique precise semantic alter-ego. We should not speak of THE semantic alter-ego, but of A semantic alter-ego that serves to give us the conditions of application, and so of truth, for the original imprecisely understood statement. For example, for the

²² “Deidealize” must always be understood as reducing, not eliminating idealization.

²³ I spell out this very compressed argument sketch in much more detail in my (to appear b).

statement, *John's height is 180cm* (understood as *close as makes no difference*), it does not matter whether we take its twin to be the statement that John's height is exactly 180cm, with height understood as a classical continuous measure, or as a discrete measure with steps of 1 millimeter in length, or as understood quantum mechanically, or... Remember, "close enough" always has to be understood contextually, close enough for what. In normal contexts the "for what" will be met with a range of theoretical choices for the precise twin statement.

Still, how, in practice, would one find a suitable precise statement that will function as the precise semantic alter-ego for a statement as normally understood? This will be part of the analysis of specific parts of natural languages. Setting up our theoretical framework does the underlying interpretive work. Finding the specific statements that do the needed theoretical work for a specific bit of language will then be an empirical question of specifying the function of the bit of language under consideration. To illustrate how this kind of consideration can go, here are some tentative examples.

Gradable adjectives that admit of absolute extremes (flat, straight, quiet...) wear suitable precise semantic alter-egos on their sleeve. We use the same term for both, and only when not clear from context need add an appropriate word or phrase to indicate which alter-ego is in question: straight exactly or close enough? Absolutely quiet or quiet enough...?

We can take the casual use of physical predicates, such as 'height' 'weight', and 'is water' to be imprecise predicates suitable precise alter-egos of which are provided by the theories of these subjects. The theories generate models that we regard as precise, where questions about how these models apply to the world are questions about their conditions of application. It is just such conditions that, in turn, bear on questions about the objective truth of statements using the original casually used terms.

A second wild card in the proposal as presented so far is the idea of functioning as an exact truth. How is this to be understood? A natural option (that need not be to the exclusion of others) is to appeal to the ways in which statements function as premises in inferences, where the notion of inference is drawn broadly, including both theoretical and practical inferences and where the inferences in question often are not explicitly given but operate in our reconstructions of human deliberation. In a world perfectly adapted to human needs, true statements would never lead to problematic conclusions. Ours is no perfect world, but the idea carries over: A statement functions as an exact truth when, or to the extent that, when we use it as a premise in theoretical or practical

inference, it tends to lead to conclusions that, when implemented in actions, objectively and in fact meet our needs and interests.²⁴

The account works smoothly when, but only when the idea of functioning as an exact truth is understood broadly, in the two following ways: First, it must include not just occurrent, but also potential cases. We say of a screwdriver that is sitting on the workbench that it functions to get screws into wood. Similarly, ‘functions as an exact truth’ has to be understood in terms of various uses, including use as a premise in inferences that might realistically come up. Second, truth is a deeply social concept. We apply the term ‘true’ to a statement when we are signaling to our community that the statement will function as an exact truth for others as well as for oneself; consequently the relevant interests have to include the interests of others in the community. These two qualifications eliminate a wide range of trivializations of the account. In particular they support the important distinction between truth and functioning as a truth. We express a real contrast when we say that a statement isn’t true, though it functions as a truth when, as is usual outside this theoretical context, we have in mind functioning as a truth for some narrower application.²⁵

9. Generalizing to more complex cases. Remember that, using the sewing drawer tape, we’ve measured John’s height at 180cm. Now I assert: John’s height is between one and three meters. There is a sense in which what I have said is highly imprecise – I’ve left open a wide interval for John’s height. But in the sense of precision in question in our discussion, this statement is, it would appear, completely precise. I’ve specified a completely precise interval and attributed a value of a completely determinate quantity, John’s height, to that interval. It would appear that in the circumstances given, the statement has no failing in precision, and certainly not in accuracy. Many have advanced such examples to me taking them to show that there is no need for any across the board alternative account of truth.

But not so fast! There is no such thing as *the* height of John. People’s heights go up and down a centimeter or so a day. If one attempted to refine to height at an instant of time (not what was intended, and in any case already an

²⁴ This way of thinking of functioning as an exact truth is suggested to me by Millgram’s thorough-going appeal to the function of inference in his characterization of “partial truth” throughout his (2009).

²⁵ A more careful statement will have to treat just what needs to be included in the range of cases for which a statement must work in order to count as true. For example an attribution of ‘short’ will count as true when it can be relied on, not for any case, but for the ones that are contextually relevant. But I should have said enough to show that the account has the resources for dealing with this issue.

idealization) there are still problems with what gets taken in: What about variation in John's posture, how much hair,...? And if these aren't enough there is always frame relativity of special relativity, problems about localization in quantum mechanics.... In fact the problem isn't just with reference to a number for the height, the height has to be specified in units which also involves some combination of idealization and imprecision. Indeed, the whole idea of spatial length involves idealization.²⁶ The statement fails to pick out any specific proposition because 'John's height' has no referent.

Perhaps we should operationalize *height*. So doing will yield a statement that is accurate but open ended, and so imprecise: Means of measurement can never be specified with complete precision.

So at the level of presupposition – presupposition that there IS a height between one and three meters - we will have reservations and options similar to those we had for *John's height is 180cm*. On the one hand we can think of this case in terms of a representation that is precise in every respect – not only is a precise interval in question, the quantity said to have a value in that interval is thought of as a completely determinate physical quantity that takes on any of a continuum of precise values. Because of the difficulties enumerated about height, and, we might add, similar difficulties with the use of the units term 'centimeters, and referring term 'John', so taken the statement does not correspond exactly to the way things are. Correspondence fails not because, as in the simpler foregoing case of *exactly 180cm* a precise statement is positively inaccurate, but because there are no such things in the world, precise heights of people,²⁷ that can be attributed a value in the interval (1m, 2m). Nonetheless, thinking of the world, or if you like modeling the world, as characterized in this idealized way functions as an exact truth in any of the circumstances in which we might be interested. So its conditions of application are always, perhaps science fiction cases aside, satisfied. This idealized understanding of the statement can be taken to be the precise member of a semantic alter-ego pair. Its imprecise alter-ego twin is the imprecise statement with an imprecise reading for 'height' (and similarly for the units, *centimeters*, and for the purported referent of 'John'.) We can take this imprecise statement to be true just in case the conditions of application for its precise semantic alter-ego are satisfied.

With hind sight, we see that the idealizations of the last section of a precise quantity of measure, units, and, in my view, also the referent of 'John' have to be treated as they have been in this section. Just as Braun and Sider insist, there is no proposition in fact picked out by the statement than John's height is 180cm. But we do devise the idealized model, a representation of part

²⁶ See my (to appear a)

²⁷ See my (to appear a) for much more detailed argument for this claim.

of a much simpler world, in which there is a completely specific referent, a quantity, height, exact units of measure, centimeters, and within which idealization we can work with the proposition that John's height is 180 cm.

Other examples will fare similarly, such as ones critics have given me over the years: There are some people in this room. Water is H₂O. There are bears in the Rocky Mountains....

The proposed analysis also enables us to see why such examples have always seemed to show conclusively that, in such cases, no qualification of "just plain true" could be required. In the (1m, 2m) interval example the interval specified is completely precise. And since the height in question is not, in the example, a borderline case, the imprecision of 'height' seems just irrelevant: In this case we know that, whatever the inaccuracy there might have been in our measurement, the height will be in (1m, 2m). But the imprecision of the term 'height' IS relevant because, borderline case or not, there is nothing in the world correctly characterized as the having of a precise height of any value, and so no proposition picked out attributing such an incompletely specified but precise height. On the other hand, leaving the idea of height open-ended leaves the statement also as not picking out any determinate proposition that might be true. This was just Braun and Sider's leading point, and in this respect the present proposal follows in their footsteps. What we here add is the idea that, while no proposition is picked out, the idealized supposition that such a proposition has been picked out functions as a truth, which functioning is interpreted as the truth of the imprecise alter-ego.²⁸

10. Consideration of four apparent difficulties. There will be a number of objections to this way of implementing "functions as an exact truth" as a basis for the literal truth of imprecise semantic alter-egos.

Where we are concerned with a statement's contribution to conclusions that are sound in practice when the statement is freely used as a premise in a range of applications, can we make sense of that statement's contribution as compared to the contributions, positive or negative, of other statements also used as premises? Clearly in practice this will usually be a hard question. But I see no reason to think that it can't be answered, in more or less detail, in specific cases. We have a closely analogous problem in discerning the contribution of parts to a complete functioning whole, a question we regard as perfectly intelligible and in many cases one we can answer. It would make as little sense to complain that there is an in principle difficulty in untangling contributions of various statements to the success of conclusions of inferences as it would be to think that there is an in principle difficulty in untangling the contribution of

²⁸ Braun and Sider cover this aspect of the analysis with a completely analyzed conception described as "approximate truth".

functioning parts to the successful operation of an automobile engine or the human body.

Some will also have the worry that *functions as an exact truth*, is a pretty imprecise notion, so that, given its role in characterizing truth, truth itself will come out as an imprecise notion. Both of two comments apply to this worry. First, in the present sketch I am appealing to an intuitive grip on the idea of functioning as an exact truth, augmented only with brief reference to role as premises in inference. But there is no reason to think that the notion cannot be made a lot more precise, most likely in a local way, that is in a way that will differ among kinds of application. Second, though there is much room for further precision, it is agreed that, as with all functional notions, *functions as truth* cannot be made perfectly precise, and so it is agreed that the notion of truth framed in terms of such an imprecise notion is itself imprecise. That is just as it should be. The precise notion drawn from referential semantics, truth of a determinate proposition, is an idealization, useful for topics such as understanding the structure of language. The present account, useful in understanding certain general questions about how language applies to the world, is framed in terms of an idea of statements themselves taken to be imprecise. If a statement is representationally attached to the world in an imprecise way, then it is natural to expect that for that statement to be true - for it to have gotten its representational content right - will likewise be imprecise.

Since our needs and interests play a central role in the account one may worry whether it provides only an irredeemably subjectivist account of truth. Not at all. Standards play a central role in the account, and we set the standards. But it is not up to us whether or when those standards are met. Whether or not standards we impose are met in a given situation is a completely objective matter.

The function of needs and standards in the account bridles truth with the question of whether statements “work”, which in turn suggests a pragmatist account of truth. Indeed, the account falls in the pragmatist tradition. But standard ways of dispatching pragmatism do not apply. The account is not in any way an epistemic account of truth. And on the present formulation familiar objections that many false statements “work” get no grip. The familiar examples appeal to statements that work for some very specific and local objective, but because false on anyone’s account fail for other objectives that do or easily might come up. This is ruled out on the present account. For an imprecise statement to be true its precise semantic alter-ego will be, strictly speaking, false according to the standards of referential semantics. But it is required to function as a truth for any interest that might reasonably come up. This corresponds to the truth of things such as having a height of 180cm, understood as close as makes no difference to the ideal; to being flat, flat enough easily to function as a table; and more generally to models that we use to get around in the world being broadly

similar enough in relevant respects to the ways things actually are, ways that are too complicated to be humanly accessible in every detail.

We can summarize what I am suggesting with a slogan that can be compared with the familiar sloganed version of pragmatism: Instead of claiming that to be true is to work we claim that to be true enough is to work well enough, where, in ways explained above, “well enough” is understood broadly. Lots of false precise statements nonetheless work well enough for the things that we care about, and when imprecise semantic alter-egos are interpreted as tailored to those needs and interests we understand their literal truth as corresponding to the ways in which the precise versions are “true enough”.^{29 30}

To repeat for emphasis something specified earlier: Working – functioning as a literal and precise truth – has to be understood very broadly. Truth has a social dimension: To characterize a statement as true is to say that others can depend³¹ on the statement in their deliberations. General trust in a statement requires that it function as an exact truth for any practical consideration that might plausibly come up for anyone in the epistemic community.³² Such certification will be correct or incorrect depending on whether the statement is true in the sense that I have proposed. The present point is that being true in this sense requires capacity to support the needs and interests of all.

11. Common ground and the illusion of exact truth. I have been speaking of expressions being precise, for example the precise member of a

²⁹ For a more detailed presentation of the compressed argument of the last two paragraphs see my (2012) section 5. Dewey scholars tell me that all of this corresponds well to Dewey’s intent.

³⁰ In note 22 to page 60 Millgram (2009) distances his account of “partial truth” from pragmatism, apparently because he is thinking of pragmatism only as an account of exact truth, in particular an account that obscures the ubiquitous utility of departing from exact truth. He writes: “[R]ecall that one well-known pragmatist slogan tells us that it’s true if it works. We have been examining cases where it works precisely because it’s not true.” One of the many differences between my account and Millgram’s is that by tying an understanding of imprecision to explicit truth, the present proposal clarifies the connection with pragmatism.

³¹ So the present analysis narrows the analogical gap between our use of ‘true’ in application to statements and the older use as in “She is a true friend”.

³² Space does not permit here addressing the worry, when historically the needs and interests of a whole community shift, can what was true become false? See my (2009, 261-3).

semantic alter-ego pair; and absence of qualification may suggest COMPLETELY precise. But I also want to claim that imprecision is ubiquitous. If I had intended COMPLETELY precise my account would undermine itself.

I don't so intend. Use of the unqualified 'precise' functions as expository convenience, simplification, idealization.³³ My usage in these respects illustrates a very general and important circumstance: The way we talk, and even more strikingly, the way we think about a subject matter, usually SEEMS to operate in terms of exact truths, unqualified in any way by either imprecision or inaccuracy. How can this be if, as I claim, there are no such truths?

I will ease this tension by appealing to what Stalnaker (2002, 701 ff.) calls "common ground". Participants in a conversation collectively presuppose a body of statements, where "[t]o presuppose something is to take it for granted, or at least to act as if one takes it for granted, as background information – as *common ground* among the participants in the conversation." 'Taking for granted' here needs to be understood as treated - always just for the purposes of the conversation – as without imprecision and without accuracy. For example, of course 'bald' is imprecise. But, if for the purposes of the conversation we agree to count someone as bald, we are there using 'bald' in the same way as we would use it if we thought it were, at least for the purposes of that application, precise. Similarly for 'accuracy'.

Stalnaker's idea of a conversational common ground applies much, much more broadly. Kuhn observes that for practitioners of a paradigm or disciplinary matrix to make any progress and to work smoothly with their colleagues they must not question the laws and other statements that are characteristic of the paradigm – they must treat them in exactly the same way as Stalnaker specifies for conversational presuppositions.³⁴ The same goes for a great deal of social deliberation and communication of information. We can put this by taking the idea of a conversation with a great deal of latitude and including a great many sustained social activities.

Why the need for conversational (in the very broad sense) presuppositions? In all these social activities we need to use statements to coordinate with our fellows. Insofar as we are paying attention to either the imprecision or inaccuracy of a statement we will be distracted from such use. We sensibly ignore such defects in statements when we believe that these defects will not interfere with pursuing the objectives in question in the activity.

³³ This was also the apparent difficulty addressed in note 17.

³⁴ See Kuhn 1970, pp. 19-20 and ch 2, 3, and 4 passim.

Once we appreciate why statements can function as presuppositions in a “conversation”, we appreciate that we can draw the contours of a conversation not just broadly, but very narrowly. Exactly the same reason for treating statements as free from any imprecision or inaccuracy that apply for conversations in the large apply in the same way for the deliberative activity of an individual. Active attention to a statement’s imprecision or inaccuracy will interfere with use of the statement in deliberation. But if the imprecision or inaccuracy of a statement are low enough, relative to present objectives, they may be ignored. Every bit as much for an individual as for participants in a conversation or a Kuhnian paradigm, efficient deliberation requires that irrelevant inexactness be ignored.

In short, whether the deliberation is social or private, we are constantly, though almost always tacitly, making judgments as to whether it is safe to use a statement as if it were “just plain true” (and in the same way as to whether a statement is “just plain false”). Our social network certifies many statements as safe to be used without qualification for anything that might reasonably come up, certifies them correctly or incorrectly depending on whether they are in fact true in the sense I have been developing. Such truths may be freely incorporated into any conversation, large or small.

Since, with rare exceptions, we treat the statements we use as true (or false) without qualification it’s not surprising that we should fall into thinking that we do have lots of unqualified truths! Statements that have been taken into the common ground are used, for the scope of the conversation or private deliberation, just as they would be if they were precise and accurate without qualification. We are sometimes aware of extra-conversational limitations of what we presuppose in a conversation or other activities, but often we let this drop out of sight. The result is an illusion of unqualified truth.^{35 36}

These considerations make clear why the proposed revision in how to think about truth requires no change in logic. Working within a framework of presuppositions and statements treated as exact, classical logic functions as accustomed. When something goes wrong we have to adjust the working framework not change our logic.

³⁵ I take the foregoing to be an important consideration, but only one, explaining the presumption that we have many unqualified truths. This prejudice runs deep in Western culture.

³⁶ In an isolated remark, Lewis (1979, 352) suggested what I have discussed in this section: “If a sentence is true enough (according to our beliefs) we are willing to assert it, assent to it without qualification, file it away among our stocks of beliefs, and so forth.”

12. Recapitulation of the status of the proposal and a forecast of repercussions. The methodology that I sketched in section 1 depicts the following status for this paper's proposal. It gets certain features of the concept of truth more accurately than the conventional idea: How do we squeeze truth out of our imperfect representations? What is it for a statement with limited accuracy and/or precision correctly to apply in our complex world? The present proposal has its own limitations. It obscures the combinatorial features of language. As explained in section 11, when we are using language as opposed to theorizing about such use it would be a disaster not to think in terms of "just plain truth". More generally, the present proposal will have to be evaluated for its own strengths and weakness, and, most importantly it is subject to refinement in BOTH precision and accuracy. The present proposal is not presented as an account that is, to the exclusion of all others, uniquely TRUE! Instead it promises to provide a productive alternative way to think about representational success in using language, a way that enables us to understand some important things particularly clearly.

Some will find my account persuasive, but only as an error theory of truth: Admittedly, and finite mathematics aside, nothing we say is true. Rather the account characterizes the application conditions under which it is felicitous to describe false statements as true. Such readers and I agree on everything but terminology. I prefer my terminology on the grounds that it seems pedantic to insist that virtually nothing we say is true really is.

The present proposal has repercussions for the projects of philosophy. For the most part it leaves day-to-day philosophical investigation untouched. We work within a "conversation" or a "Kuhnian framework", with certain starting points taken for granted. Within such a framework we proceed as we are accustomed, in the spirit of looking for a flawless account. But we shift how we think about evaluation. Rather than evaluation for the in practice unobtainable exact truth, we evaluate for strengths and weaknesses. We discard accounts with little to offer. We retain for further work accounts that have shown robust successes. We compare competing accounts. If a first largely succeeds where a second also succeeds and if the second has more defects, the first (often!) supplants the second. Sometimes two accounts may complement one another, both achieving significant but complementing success and also suffering different drawbacks. In such cases we retain both for further development. The reasonable expectation for any present intellectual horizon is that we will need more than one approach for working on different problems of one subject matter. Referential semantics and the method of semantic alter-egos, with their differential power for understanding very different aspects of language, provide a good example. For another consider the repercussions for evaluating deontological, consequentialist, and other such ethical frameworks.

For the bulk of issues in philosophy within the framework of day-to-day work it will be business as usual. But there are certain problems that are artifacts

of neglecting that traditional ways of thinking about truth involve severe idealization. For example in my (2008) I show how the Kuhnian conundrums about rational theory change are an artifact of thinking that the laws generated by a paradigm are supposed to be exact truths. In my (2004b) I show how we need to shift how we think about natural laws when we accept that generalizations within human reach will always, at best, be idealizations and that consequently laws we can state will, as applying to the actual world, always be *ceteris paribus*.³⁷ In my (2009) I begin the project of dealing with the repercussions for epistemology of the shift of understanding ‘truth’ in the formula for knowledge, “justified true belief”. For example, there has been much consideration of the contextuality of knowledge claims induced by the contextuality of justification. Parallel efforts are now required for the contextuality of truth.³⁸

³⁷ As Nancy Cartwright has argued in too many places to enumerate. This paper, together with material in (to appear b) also develop a novel and promising way of understanding *ceteris paribus* conditions.

³⁸ In *Hard Truths* (2009) Millgram is wrestling with many of the same issues addressed here but with a different model building approach. Millgram works with a notion of “partial truth” understood as an attitude towards statements (or “thoughts”) that are “inferentially restricted” (40) in the sense that one will use a statement only for a certain range of inferences: “When a thought is acknowledge to be partially true, and is nevertheless being used in inference, there is a recognized mismatch between representation and world – but not in any way that requires changing the representation.... assessing a claim as *partially* true allows for slack...*but only so much* of it” (108-9). With my modelers’ methodology I am obviously not going to urge that there is room for only one of these model building approaches, but rather that we must see, first, how well both can be developed, and then evaluate both for what they deliver by way of principled understanding of facts about language.

Graff’s (2000) (now Graff Fara) shares with me a role for interests: For example, for Graff someone counts as tall if they are significantly taller than some contextually specified norm, where significance is interest based. The account is thus only for gradable adjectives. My *functions as a truth* is likewise interest based; but my account applies to all of language and incorporates a role for interests of any kind.

Lynch’s (2011) proposes a functionalist account of truth. Explicitly modeled on psychofunctionalism, the realizers of the truth function can be very differently realized by different subject matters. Lynch’s has interesting overlap with my account. The idealized alter-egos treat, as an idealization, each predicate as signifying a property. But the connections between these idealizations and the world could work out very differently for different subject matters. The ways in which “morally wrong” is property-like could be radically different from the case of “red”. There is also room for various function-

References

Braun, David, and Sider, Theodor (2007) "Vague, So Untrue." *Nous* 41, 2: 133-156.

Elgin, Catherin (2004) "True Enough." *Philosophical Issues* 14 (1): 113–131

Giere, Ronald (2006) *Scientific Perspectives*. University of Chicago Press. Chicago.

Graff, Delia (2000) "Shifting Sands: An Interest-Relative Theory of Vagueness" *Philosophical Topics*, 45-81

Hart, H.L.A. (1961) *The Concept of Law*. Oxford University Press: Oxford

Kuhn, Thomas (1970) *The Structure of Scientific Revolutions, Second Edition*. University of Chicago Press: Chicago

Lewis, David (1979) "Scorekeeping in a Language Game" *Journal of Philosophical Logic* 8: 339-359.

Lewis, David (1983a), "Extrinsic Properties", *Philosophical Studies* 44: 197-20

Ludlow, Peter (2005) "Contextualism and the New Linguistic Turn in Epistemology" In G. Preyer and G. Peters (eds.) *Contextualism in Philosophy*. Oxford: Oxford University Press: 12-50

Ludlow, Peter (2008) "Cheap Contextualism" *Philosophical Issues*, 18: 104-129.

Millgram, Elijah (2009) *Hard Truths*. Wiley-Blackwell: Chichester, West Sussex.

Russell, B.A.W. 1923. 'Vagueness', Reprinted in Keefe and Smith, eds: *Vagueness: A Reader*. MIT Press.1997. pp 61-68

Stalnaker (2002) "Common Ground" *Linguistics and Philosophy* 25: 701-72.

realizations in "functions as an truth". Lynch's account, however, is utterly unlike mine in doing nothing to address the question of what it is for imprecise statements to be true.

Elgin (2004) has a study of how being "true enough" often does duty for what we think of as being reserved for strict truth.

Teller, Paul (2001) "Twilight of the Perfect Model Model" *Erkenntnis*: 393-415.

Teller, Paul (2004a) "How We Dapple the World". *Philosophy of Science*, 71: 425-447.

Teller, Paul (2004b) "The Law Idealization" *Philosophy of Science*, 71: 730-741.

Teller, Paul (2008) "Of Course Idealizations are Incommensurable!" in Rethinking Scientific Change and Theory Comparison: Stabilities, Ruptures, Incommensurabilities? edited by Léna Soler, Howard Sankey and Paul Hoyningen-Huene: Kluwer: 247-264.

Teller, Paul (2009) "Provisional Knowledge." In Constituting Objectivity: Transcendental Perspectives on Modern Physics. Edited by Michel Bitbol, Pierre Kerszberg, and Jean Petitot. Springer: 503-514.

Teller, Paul (2012) "Modeling, Truth, and Philosophy" *Metaphilosophy*, Volume 43, No 3. pp. 257-274.

Teller, Paul (to appear a) "Measurement Accuracy Realism"

Teller, Paul (to appear b) "Language and the Complexity of the World"

Teller Paul (to appear c) "Pan-Perspectival Realism"

Williamson, Timothy (1994) *Vagueness*. Routledge. London and New York.

Waismann, Friedrich, 1968, "Verifiability", in *Logic and Language*, Anthony Flex (ed.), Oxford: Basil Blackwell.