

1. Introduction.

You measure the temperature of a glass of water and say that the outcome is accurate – is right – to within a tenth of a degree. What does this mean? Presumably that there is some number that is, say, the temperature of the water in degrees Centigrade, and that the measurement outcome is within one tenth of a degree of that true value. The present paper will work to undermine this supposition, though at the very end I will present a way of understanding such statements that is consistent with all the difficulties that will have come before.

1.1 Restrictions.

I will restrict attention to physical quantities, though most of what I say should apply, with suitable modifications, to both the life and the social sciences. I will also restrict attention to quantities, such as mass and temperature, that can be represented with a measurement scale of real numbers, as opposed, for example, to curvature that requires a tensor. But what I discuss explicitly should apply also to such multivariable quantities.

1.2 Initial characterization of measurement accuracy. To fix on our target, we need to review some basics.

First one distinguishes between measurement indications and measurement outcomes: An indication is “what is shown on the meter”. But often such an indication can be corrected on a theoretical basis. A measurement outcome is the final result after such interpretation. Throughout I will have measurement outcomes in mind.

One can attribute accuracy to any of measurement indications, outcomes, the instruments used to produce indications, and the entire measurement system comprised by the

instrument and the theoretical basis used for interpretation. While much of what I will have to say will apply to all of these, we are best off, again, taking measurement outcomes as our primary target.

Accuracy must be distinguished from precision: The standard analogy refers to arrows shot at a target. The outcomes are accurate to the extent that they are close to the bulls-eye. They are precise to the extent to which they cluster closely together. So measurement can be extremely precise without being very accurate.²

I take the default understanding of measurement-accuracy to be what I will call “traditional measurement accuracy realism”. One supposes that there are in nature things, such as lumps of lead and glasses of water, kinds of things, such as lead and water, and quantities that pertain to things and kinds, such as mass, length, temperature, and time (pertaining to duration of processes); and one supposes that in concrete cases such quantities have values. Stated generally:

Traditional measurement accuracy realism (stated schematically for measurement of quantity, Q , with possible values, q in units, u ³, on an object or type of object, O):

Presupposition: There is in nature the quantity, Q , with value q in units, u , for object or type of object, O .

Then q' , a measurement outcome of Q in units u on O , counts as

- a) Perfectly accurate: $q' = q$
- b) Accurate (enough): the outcome, q' , is close enough to q for present purposes

c) Outcome q' is more accurate than outcome q'' : q' is closer to q than is q'' .

Accuracy understood in the traditional way is supposed to be an objective, not an epistemic matter. Realists will agree that accuracy can be estimated but not exactly known, but insist that there is nonetheless a fact of the matter, just how accurate, in the traditional sense, a given measurement outcome is.

2. Problems with traditional measurement accuracy realism.

2.1 General statement of the problem.

Traditional measurement accuracy realism fails because the terms used in the relevant statement instances fail to refer. We use terms for quantities and their values: “The temperature of the water in this glass”. Traditional measurement accuracy realism suppose that there is “in nature” some determinate quantity, temperature, or more specifically, the temperature of the water in this glass, that in this instance has some determinate value, say $20.258743\dots^{\circ}\text{C}$. My claim is that the term, ‘the temperature of the water in this glass’, does not have a referent. My reason is not in any way metaphysical. It is simply that the full facts of language use and circumstances of utterance fail to pick out any one thing to be the named quantity, temperature, or any one number to be the claimed value of the claimed quantity. While we will see a complex of detailed reasons for this failure, at bottom they are all consequences of the contingent circumstance that the world is far too complex for our language to get attached to completely determinate things, in particular, quantities and their value instances.

I must dwell on the form of my complaint because it is entirely different from what one usually hears from those known as anti-realists, and my argument will be misunderstood

if the reader falls back into thinking that I am attacking conventional realism in a familiar way.⁴ I am not claiming that there are no quantities with exact values in nature, nor, as some anti-realists would have it, that the whole idea of “things in nature” is incoherent.⁵ Indeed, there is no coherence problem in such statements because we can model what this would be like.

Rather, to repeat for emphasis, the problem is one of reference failure. Such determinate quantities as there may be fail to get attached to quantity terms, such as ‘time’, ‘mass’, ‘length’, ‘velocity’, ‘temperature’. With no determinate quantities attached to such terms, there are no determinate values for “them” to have. In addition, even if we suppose that the quantity terms do refer, we will see that determinate reference for terms purportedly referring to their values would fail anyway. We will also see difficulties with reference for terms for units, such as ‘kilogram’, ‘meter’, and ‘second’.

The problem is also not epistemic in the sense that presupposes that our terms for quantities and their values do refer, but that there are problems in knowing just what those values are. Rather the claim is failure of the presupposition, that the relevant terms have been successfully attached to determinate referents.

One immediate reaction is to say, well there are no point-valued referents, but we can always make do with an interval. But how is this interval to be understood? What one always has in mind is that the true value lies somewhere in the interval. But that takes us back to the questioned exact valued referents. In section 2.6 I will examine questions about intervals in more detail.

2.2 Reference failure source points.

There are different kinds of problems for three different kinds of what I will call “reference failure source points”. The first is comprised by quantities in the sense of a dimension as used in dimensional analysis. Mass, length, and time are usually taken as fundamental, and they figure in the characterization of other quantities, such as velocity, that has the dimensions of length divided by time. I will refer to these collectively as “dimensional quantities”. Dimensional quantities are theoretically individuated, that is identified by the role that they play in our theories.⁶

Our next reference failure source point is the units used in characterizing a quantity.

Without determinate units, no determinate quantity can have been picked out. Even if we had succeeded in specifying some quantity, say one called ‘mass’, just what quantity is in question is still open until we have said whether it is mass in kilograms, in grams, or some other unit. When traditional measurement accuracy realists postulate an independently existing value for a quantity of an object on an occasion, where objective accuracy is some measure of the difference between this and a measurement outcome, the independently existing value and the measurement outcome must be understood in terms of the same units.

Finally, I will need to distinguish between dimensional quantities and what I will call “working quantities”. Velocity is something abstract: Velocity of what? Velocity, or its absolute value speed, of sound in air is relatively speaking concrete; and speed of sound in air and speed of water in a pipe are different concretizations of the abstract, speed. One usually does not distinguish between the abstract dimensional and the, relatively speaking, concrete working quantities. In particular, metrologists appear to refer

indifferently to dimensional and working quantities as measurands, for example, VIM:
2.3⁷

Measurand: quantity intended to be measured.

But the distinction does tacitly occur in both VIM and GUM. VIM, 0.1:

[E]ven the most refined measurement cannot reduce the interval [that can reasonably be attributed to the measurand] to a single value because of the finite amount of detail in the definition of a measurand.

If one has dimensional quantities in mind, this statement puzzles because of absence of any concrete mention of refinement of “definitions” of dimensional quantities. However we see what is in question in GUM. GUM echoes VIM with:

D.1.1: The first step in making a measurement is to specify the measurand — the quantity to be measured; the measurand cannot be specified by a value but only by a description of a quantity. However, in principle, a measurand cannot be *completely* described without an infinite amount of information....

What is in question becomes clear with the following example, D.1.2:

Commonly, the definition of a measurand specifies certain physical states and conditions. EXAMPLE The velocity of sound in dry air of composition (mole fraction) $N_2 = 0.7808$, $O_2 = 0.2095$, $Ar = 0.00935$, and $CO_2 = 0.00035$ at the temperature $T = 273.15^\circ K$ and pressure $p = 101,325$ Pa.

What is the infinite amount of information here referenced? Conceivably an indefinitely long list of such potentially relevant characteristics. But more likely it is the interval left open by all such specifications. It is understood that temperature is being specified as $T = 273.15^\circ K \pm .005^\circ K$, etc.

In any case, I need the distinction between abstract dimensional and (relatively) concrete working quantities because there are vastly different problems that arise for the two. When working quantities are in question there will be some differences between type and completely concrete token cases. When discussing the speed of sound in air or the melting point of lead one has in mind characterization of the a property of a *kind* of substance – air or lead as a type. But one also needs to measure quantities for concrete instances – tokens - such as the speed of sound in the air in the Sydney Opera House at some specified time, or the temperature of the water in some specified glass at a specified time.

2.3 Difficulties with working quantities.

As relatively concrete realizations of dimensional quantities, whatever problems will arise for dimensional quantities will, ipso facto, apply as problems for their concrete realizations. But working quantities present additional difficulties. Roughly speaking, these difficulties arise in either how their dimensional abstractions are made concrete or from the fact that they are not made completely concrete. To make these additional difficulties clear, for the discussion of working quantities we will take their dimensional abstractions as given and unproblematic.

When working quantities are in question there will be some differences between type and token cases. When discussing the speed of sound in air or the melting point of lead one has in mind characterization of the a property of a kind of substance – air or lead as a type. But one also needs to measure quantities for completely concrete instances – tokens - such as the speed of sound in the air in the Sydney Opera House, or the temperature of the water in some specified glass, both at specified times.

Taking token cases first, consider a measurement of the speed of sound in the air in the Sydney Opera House at 8:00pm Jan 1, 2013. There are two difficulties. First, just what will we count as part of the Opera House? Include the vestibule? Oh, you'll protest, obvious what was intended was the auditorium of the Opera House. But to no avail. With the door open or shut? Filled with an audience or empty? Any specification of a concrete object will leave it open to some extent precisely what object is in question. Having failed to designate a determinate concrete object, there can be no determinate value that "it" actually has.

Second, the speed of sound will vary from one part of the Opera House (or the auditorium of the Opera House, or the....) to another. For example, speed of sound varies with temperature, and the temperature won't be absolutely constant throughout. There will be edge effects....

Turning to type cases for working quantities: This is the problem from VIM and GUM quoted above. The problem could be understood in two ways. First, "speed of sound in air" is open ended, as is "speed of sound in air at temperature $T = 273.15^\circ\text{K}$ ", and likewise "speed of sound in air at temperature $T = 273.15^\circ\text{K}$ and pressure $p = 101,325$ Pa". Could this list be continued indefinitely with more and more relevant features? Possibly, but that's a bit implausible, so let it pass.

But second, how are the specifications to be understood? As mentioned above, most plausibly with a temperature of $\pm .005^\circ\text{K}$ and pressure $\pm .5$ Pa; and values in the intervals will give rise to different speeds of sound.⁸ One could, on the other hand, take the specific characteristics of temperature and pressure to be intended as completely precise. But no real world sample of air has such precise values, if only because the

values would vary slightly from place to place. So at best one is talking about the speed of sound in air at... in some idealized condition, not in the real world.

2.4 Units.

The characterization of units presents a whole new raft of problems. Except for the kilogram, fundamental units are now defined using a theoretical definition. For example, currently

The second is the duration of 9,192,631,770 periods of the radiation corresponding to the transition between the two hyperfine levels of the ground state of the cesium 133 atom... This definition refers to a cesium atom at rest at a temperature of 0°K.⁹

This definition involves a number of idealizations.¹⁰ Before getting specific I need to separate out the kind of problems that will be in question for us.

To operate as a standard such an idealized theoretical definition has to be realized in some concrete piece of apparatus that will in practice function as the standard, and so doing involves deidealization from the theoretical definition. One first constructs the needed apparatus so as to minimize as far as possible the departure from the idealized definition, and one then further deidealizes using theory based adjustment of the indications physically produced by such instruments.¹¹

This need for practical deidealization in physical realization of a standard differs from the implication of idealization that we will now consider. The practical case concerns the operation of some concrete device. In examining traditional measurement accuracy realism we are concerned with, rather, whether the theoretical definition succeeds in

picking out a referent, picking out some real world characteristic, quite independently of the question of whether that characteristic can in practice be exactly realized.

The form of the problem is that the idealizations involved in a definition of a unit mean that the definition is of a unit in an idealized situation, speaking metaphorically, in a non-actual “possible world”. There is no guarantee that what is picked out for one or more such non-actual possible worlds will correspond in the way needed to any one determinate referent in the real world. Examination of cases shows that this is exactly what is in question.¹²

Let’s consider first the one unit that is still “defined” by a physical standard, the kilogram characterized in terms of the international prototype kilogram. Taking this as a perfectly precise characterization of what mass will count as a kilogram involves idealizing away variable factors, such as contaminants from the air and scratches induced when the prototype is handled in making replicas, both problems that managers struggle to minimize but can never completely eliminate. Strictly speaking, sublimation of the material of which the prototype is composed has also to be idealized away. Or, if one refrains from such idealization, there is no one mass that the prototype picks out over time because the complications such as the ones just mentioned mean that the mass of the prototype varies up and down. Even at one time there is no completely determinate real world mass that is picked out – for the same reason that gave rise to one of the complications for token cases of working quantities: Absolutely precisely, just what is, even at a fixed time, THE prototype? No one answer to this question will pick out an object that will provide the kind of standard that we assume. A policy either of including or of not including the present scratches picks out, at best, a standard that will be different

as soon as the prototype is handled. Or to give a circumstance that is utterly inconsequential in practice, strictly speaking relative motion of exactly zero is an idealization. Real world uses will involve relative motion, and so an indeterminacy in what is in question: rest or relativistic mass, and if the latter, which one? Utterly inconsequential in practice, but the realist requires a *completely precise* value.

Other standards are defined theoretically. Consider the theoretical definition of the second, just above. This definition ignores the time energy uncertainty relation that results in spectrum bandwidth. Given the bandwidth, the definition does not pick out any unique real world temporal duration. Or again, appeal to a temperature of 0°K . Nothing in the real world can be at 0°K , nor can 0°K be approached asymptotically because of the finite limit imposed by quantum vacuum fluctuations. At best the definition characterizes a temporal duration in some possible world. In fact in many possible worlds since there is no unique way in which the idealizations can be removed (What will a possible world with no quantum effects be like??) There will be no sense to be made of which of such possible worlds is “closest” to the actual world, so appeal to “closest world” won’t pick out a unique real world temporal interval.

We’re not done. The theories, general relativity and quantum field theory, used in the theoretical definition are themselves idealizations - two theories that are not unified, and of which it is at least questionable whether they are mutually consistent. These idealizations provide further reasons why, strictly speaking, the definition only gives a temporal interval in some, or really in many, possible worlds.

Definition of the meter also fails to deliver the completely determinate length that realists require. BIPM gives the definition of the meter as

The meter is the length of the path travelled by light in vacuum during a time interval of $1/299\,792\,458$ of a second.¹³

This definition inherits all the problems of the definition of the second. It involves the further idealization of the speed of light in vacuo, and the idealization of general relativity applies anew, now through its ideal treatment of distance.

2.5 Dimensional quantities.

I've saved the most vexed case for last, the case of dimensional quantities. As I mentioned, dimensional quantities are individuated by the theories in which variables for these quantities occur. But the theories in question are all idealized. So **in the real world** there are no quantities as characterized ~~in the real world~~ in our idealized theories. If they occur anywhere, it will be (again, speaking metaphorically) in the idealized possible worlds of the characterizing theories.

Take the example of mass. Is this supposed to be Newtonian mass? Relativistic mass? The mass of quantum field theory that is a renormalized quantity and so dependent on the “impact parameter” involved in its measurement? Quantum field theory is still highly idealized, so there is good reason to think that further deidealization will further recharacterize just what quantity is in question.

One wants to protest: These increasingly accurate characterizations are all of one quantity of which our theories are giving an increasingly faithful account. I will discuss the “close to” worry in a general way below. But the example of mass helps to make clear the weakness of the response. The mass of quantum field theory is so different from that of Newton that the idea that we are just refining an already very clear idea loses all plausibility. The ONLY constraint on further deidealization is that old successes be

preserved. These old successes may be preserved by radically new ideas of quantities. This can happen by the operation of a limit. In the relation between special relativity and Newtonian mechanics, one gets the latter from the former **by letting v/c go to zero**. But that doesn't make the Lorentzian metric and its geometry just a refinement of Euclidean geometry.

Let's try time: Our best theory of time is the general theory of relativity (GTR). But GTR is not quantized and current efforts to quantize GTR play havoc with the treatment of time. We don't know the outcome of this story, but at the very least there is the lively possibility that a better theory characterizing time may characterize it as differently from GTR as quantum field theory characterizes mass as compared to Newtonian or relativistic mass.

Let's try another quantity, velocity. Velocity doesn't occur as a quantity in quantum theories. When we can ignore quantum corrections one takes *speed* (magnitude of velocity) to be the limit of average speed. But the limit of averages is another idealization, one that breaks down badly even before we get to quantum corrections. And if by speed we mean an average speed, which average?

What about length? When one takes into account the indeterminateness of relative position as characterized in quantum theories, there is no such quantity. Indeed, in quantum theories length, or (relative) position, is characterized as an operator not as a real valued quantity, again, a radical departure from prior conceptions. Likewise in quantum theories momentum is a radically different kind of quantity from prior classical characterizations, like quantum mechanical position also characterized by an operator, not by a real number. These few words paper over a great many complications, but

should be enough to show that there are serious issues for the case of both position and momentum.

2.6. Repair by appeal to intervals? The realist in us all is screaming: True, no objectively occurring precise values are attached to our terms. But, objectively, suitable intervals (or other collections of values) can do the needed realist work. Here I consider this option, construed in terms of completely determinate collections of values, that is, collections for which, for each number, there is a fact of the matter whether it is in the collection or not. Later I will consider “indeterminate collections” (starting with the question of what that could even mean).

How should such an interval be understood? What one wants to say is that we are taking about an interval of values that are, in some sense, “close enough”. But close enough to what? For realism, as we have construed it, in a given problem situation there must BE a value closeness to which counts as “close enough” however that is to be understood. But for all the reasons given above, there is nothing in the problem situation that fixes the needed objective value.

For the case of working quantities there is a more careful way to make out the interval intuition. Let’s see how this goes for **speed of sound in air**. To review the problem, specifying a quantity as “speed of sound in air” is, as VIM and GUM would put it, an incomplete definition. Liquefied air? Ionized air? It is plausible that all such extreme cases can be eliminated with a short list of more specific conditions: Air at temperature $T = 273.15^\circ \text{ K}$ and pressure $p = 101,325 \text{ Pa}$. But such characterizations of the quantity are still open ended: in the present example temperature $\pm .005^\circ \text{ K}$ and pressure $\pm .5 \text{ Pa}$. The proposed solution, in the spirit of supervenience, suggests that we get our

interval by considering ALL the ways in which the characterization could be made completely precise. To put it once more metaphorically, consider the possible worlds each having some precise value for the quantities in question (temperature, pressure,), the range of possible worlds fixed by the limits in such incomplete specification of the quantities and that are otherwise maximally similar to the actual world. Our required interval (or other collection) of values will be the values in one or another of such possible worlds.

Such an interval would be objective. The statement of realist accuracy would have to be restated: Instead of distance from some one value there would have to be some relation to the interval of question. This could be done in a variety of ways, the details don't matter. This proposal collapses, in different ways, depending on how velocity is understood. Let's suppose, which is what one usually has in mind, that it is instantaneous velocity that is in question. Again, this is an idealization: there is no such thing in the real world. The proposal is to consider a range of possible worlds that differ from the real world *only* by having one or another precise value of the associated quantities, such as pressure and temperature, that are within the bounds of the interval specified in the detailed characterization of the condition of the air in which the speed of sound is in question. (For the moment we are waving the problems with both pressure and temperature, which, when reintroduced, further spoil the effort.) But with these worlds differing from the real world only by variation of the exact parameter values within the given bounds, these possible worlds will also have no instantaneous velocities. If velocity means instantaneous velocity, the proposal is empty.

The alternative is to consider some kind of average velocity in each of the relevant

possible worlds. But which? No question but that there are averages – distance covered divided by the time of travel – that will work for practical purposes. (I’m taking the appeal to “practical purposes” to paper over the problems with appeal to the distance and time of travel. This broaches problems of vagueness, to be discussed in section 4.) But the realist needs to be specific. “Pick some average that works for our current objectives” doesn’t fit the bill in the actual world, let alone in all of the various possible worlds relevant in the proposed analysis. In addition there are problems with the averages themselves. Wave or group velocity? Wave velocity is strictly defined only for a wave that extends to infinity forward and backward. And distance traveled in unit time brings in all the problems with measures of both distance and time – the problems we have already reviewed both for the units in question and for the more fundamental quantities – distance and time.

The interval intuition fails, if anything, more radically, when it comes to units. At first things look hopeful because we are told, for example, that the current practical accuracy for standards for the second is to five parts in 10^{-15} . But what does this mean? As we will learn **in more detail** below, it means that concrete standard realizations can be built to agree to five parts in 10^{-15} . It’s not yet clear what that shows about some kind of objective interval in nature. The agreement in practice clearly has some kind of controlling objective element in as much as nature makes us work very hard to get the agreement. But to what one thing “in nature”, whether point valued or precise interval, in terms of which realist-accuracy might be characterized, does this “objective element” correspond?

Unlike the case of working quantities, there is no natural candidate for the needed

interval. For working quantities one plausibly turned to all the different ways in which an incomplete specification of the working quantity might be filled in. But in addressing the idealizations involved in the characterization of a unit there is no natural or well-defined range of cases of what will count as a deidealization. The only constraint on deidealization is that past successes be preserved, that in the case of units amounts to the successes in getting real world realizations to agree at least as well as before any new deidealization. But what would be meant by the “interval of deidealizations” that might sustain the level of agreement that we now achieve in practice?

Dimensional quantities suffer, for this issue, the same problems as do units. Since dimensional quantities are abstract, unlike their concretizations in working quantities, the whole idea of an interval of refinements has no direct application. As in the case for units, any idea of an interval would have to be in terms of some range of deidealizations from the idealizations involved in the characterization of the dimensional quantity in question. It is obscure in the extreme what kind of an interval could correspond to departures of our current idealization from one or another possible “finally correct” definition of a quantity. There would have to be some kind of objective distance measure between our current idealized definition and what a “final definition” might be. As In the case of units the only current constraint on a “final definition” is that it preserve current successes. But in the case of dimensional quantities, creatures of fundamental theories, the success of a fundamental theory is entirely entangled with the work done by other theories, fundamental and non-fundamental. What would it mean to say that this success delimits some kind of “interval”, or other collection of cases reflecting facts about nature?

3. How to understand measurement accuracy

3.1 What to make of all these considerations.

For a variety of reasons, in any instance of measurement there is no completely specific value that is determined by the total situation that fixes the objective (though unknown) accuracy, in the sense of difference between **some supposed** actual value and the value that is the measurement outcome. We have considered and found wanting an effort to substitute some kind of interval or other collection of values for an objective value. Yet, there is no denying that in any actual case of measurement there is a range of values that are, as a matter of objective fact, reasonable ones that could be used, and comparison with any of which gives a measure of accuracy. Note the shift, in the last sentence, to the epistemic notion of *reasonably assigned values*. These are still objective, in as much as there is a right and a wrong, or at least a more or less reasonable, that constrain what we should do and which indirectly reflect what is going on in a world too complicated for us to know exactly.

For a sensible idea of how this works, we should look at how metrologists evaluate accuracy.¹⁴

3.2: How metrologists evaluate accuracy: Robustness accuracy. As I have been at pains to emphasize, our understanding of quantities and how they might be measured is hostage to our currently best theories. Time is characterized by GTR, temperature by thermodynamics and statistical mechanics, and so on. Also central are theories that describe the interrelation of the quantity in question with other quantities. Where time is measured by periodic motion, crucial are theories of the motions in question. The current definition of the second appeals to a spectral emission of cesium, the theory of which calls on quantum field theory. Temperature is measured by the temperature dependence

of other quantities such as volumes of gasses and liquids, the electrical properties of substances, and again spectral properties of electromagnetic emissions for heated substances. Designing and evaluating measuring instruments for temperature requires applying the theories of these substances and the relation of temperature to their other properties.

Let's look in a little more detail at how this plays out in the case of determination of units. A unit, such as the second or the kilogram, is given a theoretical or physical "definition". The theoretical characterizations require various idealizations, such as 0°K and a zero gravitational potential. The physical prototype for the kilogram functions as a fixed standard only under idealizations such as no scratches when handled and no absorption of impurities. The theoretical characterizations then must be physically realized. While the prototype for the kilogram is already physically realized, the same problems that arise for physical realization of theoretically defined units arise for a physically defined unit with the need to make copies. The physical realizations or copying depart from the idealized theoretical definitions and ideal circumstances assumed for a physical standard. To make effective use of a standard one must, in physically realizing or copying, insofar as possible, minimize these departures from the idealized definitions and conditions; and one further appeals to any relevant theory for help in further correcting for departures from the idealizations insofar as these departures still affect the physical realizations and copies. As we have seen, such deidealization cannot be done in any perfectly exact way, and what we come up with is hostage to the theories we use. Still, these theories are the best account of nature that we have, and we use them as best we can.

Metrologists work to keep track of such departures from the idealizations with what they call “uncertainty budgeting” that appeals to theory to estimate the uncertainties that arise as a result of failure to completely deidealize.¹⁵ To be sure, these departures are not from something exactly fixed in nature but from standards that are as characterized by theories that are themselves idealized. That is, it is understood that these uncertainty estimates are relative to the theories used and thus limited by the shortcomings of these theories. In consistency with all of the worries of section 2, these are not estimated departures from something fixed in nature but from the ideal depicted by what we take to be our best theories.

It is these estimated uncertainties, deployed in a robustness condition, that then provide the basis for attributing a level of accuracy to a measurement standard. In Tal’s account of the special case of the standard second one uses

two interlocking lines of inquiry: on the one hand metrologists work to increase the level of detail with which they model clocks. On the other hand, clocks are continually compared to each other in light of their most recent theoretical and statistical models. The uncertainty budget associated with a standard is then considered sufficiently detailed if and only if these two lines of inquiry yield consistent results. The upshot of this method is that the uncertainty ascribed to a standard clock is deemed adequate if and only if the outcomes of that clock converge to those of other clocks within the uncertainties ascribed to each clock by appropriate models, where appropriateness is determined by the best currently available theoretical knowledge and data-analysis methods. (Tal 2011, 1091)

We have essentially the same story for the accuracy of measuring instruments proper. One provides a theoretical model for an instrument, relying on theory to minimize, insofar as possible, the uncertainties in the sense given above. Insofar as practicable, such models will take into consideration all the factors that, according to current theory, might affect the measurement process. One then uses these models to estimate the residual uncertainties, the inaccuracies to which the instrument might still be subject, once again according to our best theories. All the estimated uncertainties are combined, and combined with the over all uncertainty in the unit standard used in the calibration of the instrument.

The estimated uncertainties, deployed in a robustness condition, then provide the basis for attributing a level of accuracy to an instrument. Tal's summary is:

Given multiple, sufficiently diverse processes that are used to measure the same quantity, the uncertainties ascribed to their outcomes are adequate if and only if

- (i) discrepancies among measurement outcomes fall within their ascribed uncertainties; and
- (ii) the ascribed uncertainties are derived from appropriate [as described above] models of each measurement process. (Tal 2012, 175)

Uncertainties that satisfy this robustness condition qualify as reliable measures of the accuracies of the measurement outcomes of the instruments in question.

3.3 But why should such uncertainties count as measures of accuracy? One may take the robustness condition to proceed in the following spirit.¹⁶ The world is too complicated for us to be able to describe it exactly as it is. We have to rely on a network

of (not always exactly consistent) idealized theoretical accounts. But we use these accounts precisely because they give us a good enough picture to get along for a wide range of objectives. The uncertainties that figure in the robustness condition are not interpreted as uncertainties of departure from the realists' actually occurring values¹⁷ but as departures from values that we can suppose would occur in the idealized circumstances described by our theories. Broadly, our composite idealized accounts are good enough to be highly reliable and it is just a special case of this over all reliability that we won't get into trouble by treating departures from supposed idealized values of idealized quantities characterized in idealized units as departures from postulated actually occurring values of real quantities described in exactly characterized units. In the larger idealized picture of the subject matter the measurement outcome is off by some (not exactly known) definite value from WHAT IT WOULD BE IN THE (OR SOME) SIMPLIFIED WORLD characterized by our idealized larger picture.¹⁸

Accuracy realism fails because of reference failure, and reference fails because of, a fact that we too easily let drop out of view, the ubiquitous idealizations of our theoretical accounts of the world. We forget the idealized status of our theories precisely because they work so well and so broadly. Generally speaking, we get on successfully treating the world as characterized by the idealized dimensional quantities, specified in idealized units, and then applied more specifically with the idealized concrete versions provided by working quantities. In short, we proceed AS IF the presupposition of traditional measurement accuracy realism were true. In other words, the presupposition of accuracy realism is itself an additional idealization, or perhaps a collective application of prior idealizations.

Measurement standards function for us as the bench marks against which measurement accuracy is evaluated. But that comes down to saying that we treat objects as having values for quantities as characterized in terms of our current measurement standards. On the one hand, we know that these standards are always susceptible to improvement, in ways in part marked by the ascribed uncertainties. But at any moment we can do no better than to treat the world as characterized in terms of these standards, that is, as if the world were just as so characterized. Acknowledging that improvement is always an option comes to acknowledging that using a standard as our guide to the world is an idealization.

The robustness condition is essential to the success of so proceeding. The condition functions as a prescription to check, check with great thoroughness, that the various ways in which we assign values to quantities as described in our theories all fit together well enough not to engender difficulties.¹⁹ The robustness condition functions precisely to insure that taking the presupposition of accuracy realism, made concrete in terms of our measurement standards, as an additional idealization or collective application of antecedent idealizations, does not spoil the larger operation of the sketch of the world provided by concrete application of our interconnected idealized models and theories.

The current proposal is not to scrap the concept of traditional accuracy realism in favor of some substantially different concept. Rather I am urging a change in how we think about the concept. We apply the familiar concept but no longer in a traditional realist spirit.

Instead we appreciate its status as an idealization. Consider some specific measurement situation with an object of measurement being evaluated for the value of some quantity as characterized by our relevant current theories. Satisfaction of the robustness condition

insures that if one were to use any realization of the available measurement standards for this quantity one would get the same value up to the tolerances characterized by the uncertainty budgeting. Given this reliable consistency one won't get into trouble by idealizing, by thinking of the situation as one in which there IS a quantity characterized by our theories, a unit set by the measurement standards, and that the object has a value for that quantity in those units. This last is just to say in the material mode exactly what is reexpressed in the formal mode by saying that the expression

There is a quantity characterized by our theories, a unit set by the measurement standards, and that the object has a value for that quantity in those units

has precise referents, the quantity, the units and the value in question. We know that these expressions do not have referents, but there is much practical advantage and no harm is done by treating them as if they did.

As for the accuracy of some instrument that is not part of the system of measurement standards, we think of it, within the scope of the idealization, as the difference between the measurement result and the supposed actual value. Of course even if the world were as in the idealization, the best we could do to get that supposed value would be the values of one or another measurement standard, qualified by the uncertainty budget. But since in the real world the very best we could presently do would be exactly those results of one or another measurement standard, in practice what we can have is exactly what we would have if the world were as in the idealization. With the robustness condition in place the idealization can't get us into trouble.

4. Quantities and units understood as vague.

Some readers will have been thinking throughout: The problems here are all problems having to do with vagueness. So there is no special problem here, nothing problematic over and above whatever general problems there may be with vagueness.

I agree, at least for the case of working quantities and units. The proposal, however does not work for the case of dimensional quantities. In this section I will examine the connection and argue that treatment in terms of vagueness and in terms of idealization as in the last section are really two different ways of getting at the same thing. Certain advantages accruing to working with idealization, starting with the circumstance that framing in terms of idealization, gives a uniform treatment of dimensional quantities along with units and working quantities, also thereby providing a kind of generalization of the notion of vagueness.

On its face, ‘accurate’ is vague in exactly the way that ‘flat’ is vague – ‘accurate’ is not the target there. Rather it is the expressions that have the form of picking out units, quantities and their values.

Compare:

The temperature of the water in this glass

The time at which John arrived home.

There is no one temperature that counts as *the* temperature of the water in this glass. If you think that temperature is an intrinsic quality of objects no one number will do – any real body of water in a glass will have some temperature gradients. Perhaps you want to take a statistical mechanical definition of temperature, the mean kinetic energy of all the molecules in the glass – but this too will suffer fluctuations and is, in any case, a classical idealization. Likewise there is no one precise moment that counts as *the* moment at

which John arrived home. When he pulled his car into the driveway (and just which moment was that)? When he stepped over the threshold? When he hung his hat on the hat rack...? ‘The time at which John arrived home’ is vague, and in an analogous way so is ‘The temperature of the water in this glass’.

Note that this proposal importantly differs from the interval proposal considered in section 2.6: There is no determinate interval of values that could count as the time of John’s arrival; and likewise no determinate interval of values that could count as the water’s temperature. Below we will consider how to make sense of a contrasting “indeterminate” collection of values.

If working quantity terms are vague, then there appears to be a simple way to characterize ‘accurate’. Let’s call the imprecisely characterized collection of values that would work for the temperature of the water in this glass the “temperature value collection”. We could then characterize

Accurate (enough): Close enough for present purposes to any one (or almost any one) of the values in the temperature value collection.

Problem solved! On this analysis, ‘accurate’ comes out as doubly vague: vague in the “enough” (compare: flat (enough)), but also vague in the “the values in the in the temperature value collection”.

We understand this approach precisely as well as we understand “temperature value collection.” Again, section 2.6 rejected the option of saying that, though imprecisely characterized, there must BE some determinate interval or other collection of values that is in question. We need to develop some alternative way of thinking about what is going on when we talk about such intervals or collections.

Let's work this through for "the time at which John arrived home". Suppose that we check the security camera and find that John turned off his car at 4:59:32, crossed the threshold at 5:00:02 and his hat hit the hook on the hat rack at 5:00:05. For virtually any practical purpose that might come up you could use any of these times for "the time at which John arrived home", as well as many others that are, from a practical point of view "close enough", though 5:00:00 would be the obvious practical choice. Which numbers could be used is open ended in the sense that which ones would be appropriate choices depends on what is at issue, in turn fixed by the context, but "the context" itself will shift from case to case and in no case will be specific in every respect. At the margins one is free to cut the edges as one likes, and when the margins are fine enough choices will be arbitrary.

There is no determinate collection of values that qualify. There are only the practical questions of what numbers will serve, and how well, for practical issues.

Yes: the approach illustrated in this example supplies an approach that could be applied very generally to the phenomenon of vagueness. Tal's robustness analysis is attractive because it provides a basis for making out this kind of thinking in an exceptionally general and coherent way for the case of measurement accuracy.

I do not know of any explicit development of this approach to vagueness in the vast vagueness literature. The usual way of dealing with the worry of indeterminately specified collections is the hierarchy of higher order borderline cases. Such developments may provide interesting formal constructions, but they are terrible models of vagueness of terms in natural languages and in particular in the languages of science. A borderline case is a case in which one is appropriately unsure about what to say. A borderline-

borderline case would be one in which one is appropriately not sure whether one is appropriately unsure about what to say – something that in some cases we can make sense of, but something that in practice arises extremely rarely, if ever. The third order case goes beyond any normal human capacity or need, and so beyond anything that corresponds to the function of human language. The pragmatic approach to vagueness fits these conditions perfectly.

Understanding vagueness as a practical question of applicability makes it easy to see the connection with idealization. In the case of the temperature of the water in this glass, given practical and theoretical questions of applicability that might come up we have some latitude as to which number to use for the temperature. Choice of any one is an idealized description of a much more messy real world situation. An idealization is, strictly speaking, false. But within its domain of applicability one can use it as if the world were just as the idealization says it is. That is, for a suitable range of practical or theoretical questions the idealization functions as a precisification of a corresponding imprecisely characterized situation. In the case of the temperature of the water in our glass, postulating a precise value for the temperature is a precisification of the imprecisely characterized temperature value collection, one that is appropriate just when it is one of the values that arise in the practical analysis of the kind we have seen above. I call precisely stated idealizations and corresponding imprecise, or vague, characterizations “semantic alter egos” because they are different ways of accomplishing the same semantic work.^{20 21}

But working with idealizations has the advantage that it will apply in the treatment of dimensional quantities where, I will now argue, thinking in terms of vagueness no longer applies.

To understand a vague term requires understanding how to make it more precise. This we can easily do for working quantities. For units I could make a case either way, depending on how we make more precise the vague “understanding how to make it more precise”. But dimensional quantities can’t be forced into this mold for the kinds of reasons that already came up when we discussed the precise interval option for dimensional quantities. What would count as precisifications for terms for dimensional quantities would be de-idealizations. But for our currently most detailed theoretical account of a dimensional quantity we have no idea how to de-idealize – if we did we would already have these proposed theories on the table!

To make this out in more detail requires addressing a complication. If the issue is put, not in terms of an attribute of theory, idealization, but instead an attribute of language, vagueness, we have also to bring in the phenomenon of ambiguity because many of the relevant idealizations in question will correspond to ambiguity rather than vagueness. Vagueness – susceptibility to precisification from an indefinite range of refinements – and ambiguity – susceptibility to disambiguation from a determinate, very limited collection of determinate meanings – are not the same phenomenon.²² For many theoretical considerations about language they must be distinguished. But for our purposes we can lump them together. The two share the relevant feature that to understand a vague/ambiguous term in a way that involves awareness of the vagueness/ambiguity requires knowing how to make the term more precise/knowing how

to disambiguate the term. While not absolutely clear, it is at least odd and/or misleading to say that a term, as used in a language community, *is* vague or ambiguous even though no one in the community has any awareness of that vagueness or ambiguity or any understanding of how to precisify or disambiguate the term.

The kind of problem we are considering for dimensional quantities turns on variation in the discrete parameter, theory. Consider, for example, mass. As the term is now used it is ambiguous, between rest mass and relativistic mass. But this is only post 1905! I submit that as used before 1905 the term ‘mass’ was not ambiguous – it referred to Newtonian mass, now best understood as rest mass. Before 1905 no one knew of the relativistic alternative so disambiguation was not an option.²³

The case for dimensional quantities differs from that of units and working quantities in two ways. First, it is ambiguity, not vagueness that is in question. From the perspective of our present interests this is an irrelevant difference. But second, when it comes to the dimensional quantities as characterized in our currently most detailed theory, the terms for dimensional quantities do not count as ambiguous. As noted above, if we could disambiguate this would be by appeal to more detailed theory that, in the cases in question, we do not have.

But the cases in which we can precisify/disambiguate and those in which we can’t still have in common the underlying source – idealization. We can appreciate that our characterization involves idealization. But we may or may not know how, at least to some extent, to de-idealize. When we do know, we have vagueness or ambiguity; when we do not know we don’t. This is the reason for which our present subject is more

perspicuously approached in terms of idealization rather than vagueness and ambiguity.

5. Concluding Thoughts.

I have argued for the systematic failure of reference for referring terms for quantities and their values. We have seen this failure as a kind of generalized kind of vagueness (and ambiguity). Since vagueness is a ubiquitous aspect of language, in and out of science, this suggests that reference failure is likewise a very general feature of language. Braun and Sider claim just this, taking this circumstance to be sufficiently obvious that no argument is required:

[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] (Braun and Sider 2007, 134)

We can see section 2 as showing in detail that this is so in the special case of terms for quantities, their units, and their values.

Just as section 2’s problem of reference failure – aka generalized vagueness/ambiguity – generalizes to all human representation, I urge that the response in section 3 likewise generalizes: It is through idealizations that we know the world. The world is too complex for us to have representations that characterize it exactly, that is with both perfect precision and perfect accuracy. Our representations always fall short in one or both of these two ways. This is as true of perceptual as of theoretical knowledge. But we do know a great deal. Knowing the world is knowing the world through idealizations, and insofar imperfectly.

It is a fair question, what is it to know the world through idealizations? This is a question

on which I have touched in many other articles²⁴ but which needs much more thought and discussion. Indeed, it requires a wholesale overhaul of our understanding of human knowledge. For the moment I will leave it with the suggestion that the present treatment of measurement accuracy and its appeal to Tal's robustness condition provides an exemplar that can usefully guide our thinking.

The problem I have discussed in section 2 is the semantic problem of reference failure, not an epistemic problem of difficulty in knowing values that are alleged to have been fixed. But the suggestion of this section is that ubiquitous reference failure gives rise to a very different epistemic limitation, that we know the world only through idealizations. Philosophical tradition to the contrary notwithstanding, knowing imperfectly is still knowing what the world is like; in particular, that **the world** is very like one occupied by such and such idealized objects with such and such idealized characteristics. One doesn't have to get it exactly right about what things there are and their properties. There is a difference between getting things wrong in ways or to an extent that do not presently matter and getting things badly or completely wrong. Complete precision and accuracy is not humanly attainable and also not needed. Imperfect knowledge is still knowledge of the world, we can add redundantly, of the way the world is REALLY. This isn't traditional realism, but it is the sensible way in which we should have been understanding realism all along.

Acknowledgement. As readers will see, this paper draws heavily from Eran Tal's (2011 and 2012). In effect, I am offering my own more detailed interpretation of parts of his work. I have also profited greatly from comments on a draft of this paper from Tal, as

well as from Michela Massimi, Bas van Fraassen, and Ron Giere.

ENDNOTES

¹ As readers will see, this paper draws heavily from Eran Tal's (2011 and 2012). In effect, I am offering my own more detailed interpretation of parts of his work. I have also profited greatly from comments on a draft of this paper from Tal, as well as from Michela Massimi, Bas van Fraassen, and Ron Giere.

² I have written about measurement precision in (Teller 2012)

³ Reference to units is a short way to cover the point that what is postulated are not values as numbers, in some Platonic sense, but a ratio or other relation between the quantity Q , as it applies to O , and the quantity Q as it applies to some reference object (e.g., international prototype kilogram) or condition (e.g., the radiation spectrum of cesium) that sets the units.

⁴ I do not consider the view I present in this paper to be ant-realist. Indeed, as I will explain in the paper's last paragraph, the present view is the sensible way that realism should have been understood all along.

⁵ Nor would I claim that there *are* such things "in nature", whatever that might mean. This paper is entirely agnostic about this question. If the reader **MUST** know what my private view is on this matter, let me just say that it is deeply Kantian.

⁶ As argued by Tal who concludes that

In order to individuate quantities across measuring procedures, one has to determine whether the outcomes of different procedures can be *consistently modeled in terms of the same parameter in the background theory*. If the answer

is ‘yes’’, then these procedures measure the same quantity *relative to those models*. (2012, 84 Italics in original)

This quotation also makes it clear that Tal is here referring to dimensional quantities, as opposed to what I will below call working quantities.

⁷ I will be referring to two documents published by the Joint Committee for Guidelines in Metrology (JCGM): the *International Vocabulary of Metrology*, or VIM, and *the Evaluation of Measurement Data – Guide to the Expression of Uncertainty in Measurement*, or GUM. The references will be by section number.

⁸ If, with VIM and GUM, we take “definitions” of quantities to include such detailed specification of properties, and with these differing quantity values count as different quantities. On this interpretation no specific quantity has been picked out.

⁹ Bureau International des Poids and Measures,
http://www.bipm.org/en/si/si_brochure/chapter2/2-1/second.html

¹⁰ I will use ‘idealization’ very broadly to characterize any representational inaccuracy. This can, but need not be understood as in comparison with some absolute standard. The alternative is to think of inaccuracy of a representation as what is so characterized from the point of view of some other representation that improves on the first in the sense of preserving all past and improving on the descriptive successes of the first representation. (Clearly in this note I am using ‘accurate’ and ‘inaccurate’ in a much broader sense than in the rest of the paper.)

¹¹ See Tal (2011, 1088-1090)

¹² It is a sensible question whether the problems that use of idealizations generates here are special to the case of measurement, quantities, units, and accuracy; or whether they

are really more general. I urge the reader to put this question aside for the moment and focus on the arguments. In section 5 I will address the question of whether these problems are really more general.

¹³ BIPM, http://www.bipm.org/en/si/base_units/metre.html

¹⁴ In what follows I am closely following Tal (2011) and (2012 chapter 4).

¹⁵ I follow Tal (2011) especially pp. 1090-1093

¹⁶ Tal appears to have a different attitude towards his robustness condition. He never considers possible failure of the presupposition of what I am calling traditional measurement accuracy realism. See his (2011, 1094).

¹⁷ Although they might also be that. Remember that this paper does not argue that there are no potential referents in nature but that, should there be such, our terms fail to attach to them.

¹⁸ Remember that I am using ‘idealization’ very broadly to characterize any representation that stands to be improved by being made more accurate in some way where “improved” may be understood as relative to some descriptively more successful representation. See note 10.

¹⁹ Note that this way of thinking of the robustness condition differs from thinking of it as a regulative ideal of bringing language and reality into perfect alignment. I don’t think that these attitudes exclude one another. It would be a very useful way to further explore these issues to consider the pros and cons of both these attitudes.

²⁰ See my (2011 and forthcoming) for much more detail.

²¹ Once the connection between idealization and vagueness is made, section 2 functions as an extended argument against epistemic accounts of vagueness.

²² See Bromberger (2012, 75-8 and passim.)

²³ A slightly more careful version of this account: Prerelativistically, ‘mass’ was ambiguous between inertial mass, gravitational mass, and the pretheoretic quantity of matter. Newtonian inertial mass was non-relational, and in special relativity it is replaced by relativistic inertial mass that is relative to an inertial frame. This makes inertial mass, post 1905, ambiguous in a way similar to the ambiguity in ‘heaviness’, ambiguous between the relational quantity, weight, and the non-relational quantity, mass, a dimension of relationality that layers on the foregoing. I differ from van Fraassen (2002, 115-6) who at least suggests that ‘inertial mass’ was somehow tacitly ambiguous before 1905.

²⁴ See my 2004, 2009a, 2009b, 2001, and forthcoming.