

Semantics as Measurement

Derek Ball
University of St Andrews

One kind of project in natural language semantics aims to explain various features of sentences (or utterances, or discourses) – including facts about entailment, acceptability, interpretability, and truth (in various sorts of scenario)¹ – in something like the following way: entities (including especially set theoretic entities such as functions) are assigned as the *semantic values* of simple expressions; rules are set out that determine how the semantic values of complex expressions are determined by the semantic values of their constituents and their structure;² and the semantic values of complex expressions (especially sentences) are associated with predictions about the facts to be explained. For lack of a better term, call this sort of view *formal semantics* (without pretence that no other approach is worthy of being so called).³

Formal semantics is not lacking in detractors (Lepore, 1983; Lepore and Ludwig, 2007). But despite the relative youth of the project, there has been very great progress: we now have sophisticated explanations of phenomena that were not even known to exist 40 years ago. Though there are very many unanswered questions, it seems hard to imagine that any unbiased observer could reckon formal semantics conceived in something like the way I have described as other than a progressive research programme. But what, exactly, are we doing when we do semantics? To begin with the very basics:

¹ See Yalcin (2014), Larson and Segal (1995, ch. 1) for discussion of the sorts of phenomena explained by semantic theories.

² The usual rule (and what I'll assume in the examples below) is function application, according to which the semantic value of a complex whose constituents have as their semantic values a function and an object in its domain is the result of applying the function to the object (see e.g. Heim and Kratzer (1998, ch.2 and *passim*)); though see Pietroski (2005) for an alternative. The usual assumption is that the structure in question is a level of syntactic representation known as Logical Form or LF, though this too is controversial (see Jacobson (2002) for discussion). The view sketched in this paper may go some way towards defusing these controversies.

³ This sort of approach is often called *model theoretic semantics*, but this term is potentially misleading, since it is debatable whether model theory per se (as opposed to the use of set theoretic entities such as functions) plays any substantial role. *Montagovian semantics* would be another possibility, but contemporary approaches eschew many crucial features of Montague's work.

we are trying to characterise some phenomenon – a phenomenon we might pre-theoretically have thought of as *meaning*. But there is little agreement about the nature of this phenomenon: some hold that it is a feature of a bit of the human brain, others a social convention, others a bit of set theory. And even if we agreed on what we are characterising, there are difficult (and rarely discussed) issues about *how* we are characterising it. Is semantic theorising a matter of stating facts about meaning? Or perhaps, instead of stating some facts, we are producing a sort of a model of the phenomenon, or spinning a useful fiction (as proponents of certain anti-realist views of science might maintain)?

Call a theory that answers these questions about what and how a semantic theory represents *an account of the nature of semantic theory*. The aim of this paper is to lay the groundwork for such an account. We begin with an apparent problem for at least some versions of the formal semantic project: *the problem of radical contextualism*. It is often taken for granted that the semantic values of sentences must be characterised in terms of truth: they are truth values, or functions from worlds to truth values, or something similar. (So decreed David Lewis: “Semantics with no treatment of truth conditions is not semantics” (1970, p. 18).) But truth conditions seem to vary with features of context. And though semanticists have familiar recipes for treating certain simple sorts of context sensitivity – for example, those exhibited by pronouns like “I” (Kaplan, 1977) – the features of context that can be relevant seem too many and too varied to be treated in this way. The case for this claim is made by examples: Pia’s utterance of “The leaves are green” (said of the green-painted leaves of a naturally red-leafed Japanese maple) can seem true if made to a photographer looking for a subject for a photo-shoot, but false if made to a botanist looking for material for an experiment in green-leaf chemistry (Travis, 2008, p. 111-2).

It cannot simply be assumed (or even decreed) that formal semantics must be truth conditional: there are versions of formal semantics – for example, dynamic views, according to which semantic values represent the way utterances change the contexts in which they are made – that do not deal in truth conditions. But these views do not escape the radical contextualists’ objection. Even the dynamic semanticist, unless her explanatory ambitions are extremely limited (say, to giving an account of certain sorts of anaphora), may well think that the two utterances of “The leaves are green” change the contexts in which they are made differently, and that these differences should be represented (somehow) by her theory. Moreover, typical formal semantic theories use natural language expressions in the metalanguage; for example, it might be claimed that the semantic value of “green” is

that function that maps an object to a certain value just in case *it is green*, and to some other value otherwise. But if “green” is context sensitive, then any such theory will either be linked to some particular context (and so presumably will deliver the wrong result if applied in other contexts); or the context in which the formal semantic theory is stated will be so impoverished as to fail to resolve the context-sensitive features of “green” (so that our specification of the semantic value of green will fail to pick out a function, in much the way a theory that said that the semantic value of “that” is that would fail to express a claim if stated in a context where no object is relevantly demonstrated or salient.) So the formal semanticist does not avoid the problem merely by avoiding truth conditions.

What responses are available? There are at least three:

Strategy 1: Explanation One can offer an account of how the truth conditions of the relevant adjectives are contextually determined. For example, Szabó maintains that “An object is green if some contextually specifiable (and presumably sufficiently large) part of it is green,” (2001, p. 138) and hence that the semantic value of “green” maps objects to truth values only when supplemented by a contextually supplied specification of a part of an object. An utterance of “The leaves are green” might require the leaves to be green only on the surface (when speaking to the photographer), or green on the inside too (when speaking to the botanist).

Strategy 2: Denial One can maintain that the semantic value of “The leaves are green” is the same in both of the described circumstances,⁴ and so in particular, if semantic value is or determines truth values, then the truth values must be the same (e.g., Cappelen and Lepore (2005); Borg (2004)).

Strategy 3: Retreat One can give up on formal semantics. Retreat can be complete (involving the complete abandonment of semantic theorising, at least as it is traditionally practiced), or only partial (retreating to the claim that, though semantics has some role to play, that role is quite limited).

Each of these strategies has its fans, but each seems subject to serious objections. Strategy 3 in its complete version seems rashly to abandon a progressive research programme with significant achievements, and even par-

⁴Or at least, this is so if we ignore factors (such as tense) that seem irrelevant to the present case.

tial versions can seem to give up on a still-developing programme too soon.⁵ Strategy 2 seems simply to shift the burden to a (typically under-described) pragmatics or speech-act theory, leaving the role and significance of semantics unclear. And though strategy 1 may seem promising when applied to any particular case, it could only work if it can overcome the general problems associated with using context-sensitive vocabulary in the metalanguage, and in any case it seems unpromising as a general response to the phenomenon: there are too many possible variations in context that can make a difference, so many that no single solution of this kind (or list of reasonable length) can hope to explain them all. (Consider, for example, “The leaves are green” said of leaves that have been (not just painted on the surface) but saturated through-and-through with artificial green dye; or for that matter, “That is green” said of the ink that looks black in the bottle but green on paper, or of the t-shirt that looks green, but only under UV light.)

Nonetheless, my aim in this paper is to defend a view of semantics that motivates a version of strategy 1. My plan is to begin with a case in which it is extremely plausible that a semantic theory does real explanatory work, and then to develop a theory of the nature of semantic theorising that puts us in a position to understand semantic explanation of this kind. On the view that emerges, semantic theorising has much in common with developing a scale of measurement for an ill-understood phenomenon. Attention to cases from the history of science in which a scale of measurement was developed, and to the practical knowledge necessary to understand measurement via the mathematical resources of measurement theory, will put us in a position to see why strategy 1 is right.

The bigger picture is that the question of *how semantic theories characterise their targets* matters. There are many kinds of theoretical activity across the sciences. Understanding semantics requires understanding the representational role that our semantic theories play; as we will see, this role is as much a matter of stipulatively setting out representational conventions as of stating facts.

1 Gradable Adjectives

So far, I have gestured at one approach to semantics and claimed that it is a sterling example of progress, with many successes to its name; but I have

⁵It should be noted, though, that some versions of this strategy (for example, relevance theory) are progressive research programmes in their own right. Nothing in this paper tells against these programmes, except insofar as they regard themselves as motivated by the idea that formal semantics is demolished by radical contextualist arguments.

not actually described any such successes. In this section, I want to sketch one example of a formal semantic theory that is capable of doing interesting explanatory work. This theory is controversial, and I do not purport to be giving anything like a full defence (or even a fully explicit description) of it; instead, the goal is to give a concrete example of explanation in semantics, so as to be able to draw some conclusions about the structure of formal semantic explanation.⁶

Consider gradable adjectives – adjectives like “hot”, “rich”, “tall”, “beautiful” – that intuitively pick out features that can be had to a greater or lesser extent, typically can combine with degree modifiers (“very”, “quite”), and have comparative forms (“hotter”, “more beautiful”). It is very plausible that such adjectives are context sensitive;⁷ what counts as “hot” will vary depending on whether we are talking about the inside of a house, the inside of an oven, or the surface of a star.

A natural first suggestion would have it that gradable adjectives pick out relations to comparison classes, which are typically supplied by context. For example, to say “Mary is tall” in a context where philosophers are salient, or to say that Mary is tall for a philosopher, is to say that a certain relation holds between Mary and the class of all philosophers – roughly, the relation of being significantly taller than the average for the class. This explains the context-sensitivity of “tall”, since different contexts will supply different comparison classes. But there is a great deal of data that it does not explain.⁸ For example:

Entailment Facts Gradable adjectives often come in pairs that intuitively attribute the same sort of feature, but do so from different perspectives: “tall” attributes height in a “positive” way, “short” in a negative way,

⁶I learned late in the process of writing this paper (from King (2014)) that Michael Glanzberg, in unpublished work, uses the same example – the semantics of gradable adjectives – to motivate some related conclusions. Glanzberg focuses on the metaseantics of contextual parameters. His claim is that certain contextual parameters must have an “indirect” metaseantics – i.e., that their values are not determined directly by speaker intentions or easily accessible features of the situation of utterance. This is a special case of the point I will go on to make in the next section: on the semantics to be considered, it is puzzling not only how the values of the relevant contextual parameters are determined, but also how the context-invariant semantic values of gradable adjectives are determined. On the account I develop, theorists’ stipulations play a significant role in fixing both the values of contextual parameters, and the standing semantic values.

⁷Pace Cappelen and Lepore (2005).

⁸I am focusing on some simple examples (following Kennedy’s discussion in his (1997; 1999; 2001), but there are others (which can also be accommodated in Kennedy’s semantics), including the distribution of degree modifiers (Kennedy and McNally, 2005).

“rich” attributes wealth in a positive way, “poor” in a negative way, etc. There are systematic entailment relations between sentences involving different members of a pair, particularly in their comparative forms. For example:

- (1) If Jane is taller than John, then John is shorter than Jane.
- (2) If climbing volcanoes is more dangerous than hang gliding, then hang gliding is safer than climbing volcanoes.

Cross-Polar Anomaly It makes good sense to say:

- (3) Mary is taller than John is.

Given standard syntactic assumptions, this is equivalent to something like:

- (4) Mary is taller than John is tall.

By contrast, (1) and (2) are anomalous.⁹

1. ? Mary is taller than John is short.
2. ? Mary is slower than John is fast.

Cross-modal anomaly Similarly, it is very difficult to make sense of (1) and (2):

1. ? Mary is smarter than John is tall.
2. ? Mary is faster than John is beautiful.

Kennedy (1997; 1999; 2001) proposes that all of these data can be explained on the hypothesis that the semantic value of a gradable adjective is a measure function: a function from an individual to a degree on a scale,

⁹ Or at least, they are anomalous on one very natural way of understanding them. It is possible to generate contexts in which these sorts of sentences sound okay. On Kennedy’s view, the acceptable readings of these sentences do not directly compare (say) Mary’s height with John’s height (or Mary’s positive degree of height with John’s negative degree of height), but the extent to which Mary’s positive degree of height deviates from some norm of tallness and the extent to which John’s negative degree of height deviates from a norm of shortness (Kennedy, 1997, ch. 3.2). These readings are marked in ways that make them fairly easy to set aside; for example, the acceptable reading of (1) entails that Mary is tall and that John is short, while the anomalous reading entails neither of these (just as (3) does not entail that Mary is tall). The same observations can be made about the examples of cross-modal anomaly, to be discussed below.

where a scale is understood to be a linearly ordered set of points, and a degree is an interval on the scale. It is perhaps easiest to understand Kennedy's semantics by beginning with his treatment of comparatives:

- (5) Derek is taller than Brian is.
- (6) Barack is more beautiful than Michelle.

Abstracting away from syntactic details, Kennedy's idea is that the semantic values of *tall* and *beautiful* are functions from individuals to degrees, and the semantic value of names like *Derek* and *Michelle* are individuals, and that these functions can be applied to these individuals. The semantic value of the comparative suffix *-er* (or of *more*) is then taken to be a function from degrees to functions from degrees to the semantic values of sentences:

$$(7) \quad \llbracket \text{-er} \rrbracket = \llbracket \text{more} \rrbracket = \lambda d_1 \lambda d_2 : d_1 \subseteq d_2 \text{ or } d_2 \subseteq d_1 . d_1 \subset d_2$$

In other words, $\llbracket \text{-er} \rrbracket$ maps two degrees to Truth just in case the first is a subset of the second, and is undefined just neither is a subset (or improper subset) of the either. A sentence like (5) will be mapped to Truth just in case the degree of height that $\llbracket \text{tall} \rrbracket$ maps Brian to is a subset of the degree of height that $\llbracket \text{tall} \rrbracket$ maps Derek to – that is, if Derek's degree of height is greater than Brian's.

Non-comparative occurrences of gradable adjectives are analysed as implicitly comparative.

- (8) Mary is tall.

Simplifying somewhat, a plausible treatment has sentences like (8) contain an unpronounced element, the semantic value of which is a function that maps a degree (the result of applying $\llbracket \text{tall} \rrbracket$ to Mary) to 1 just in case that degree is a superset of a contextually supplied degree – that is, just in case Mary's degree of height is greater than a contextually determined norm of tallness.

In order to explain the entailments between sentences involving positive and negative adjectives, Kennedy appeals to a hypothesis about the ontology of degrees – in particular, the relations between the ranges of the semantic values of the members of antonymous pairs like *tall* and *short*. For any given scale and any given object, the object's positive degree of the quantity relevant to that scale is the complement with respect to the scale of the object's negative degree. For example, if my degree of tallness is the interval on the scale of height from 0 to some point p , then my degree of shortness begins at p and occupies the rest of the scale. This ensures that if someone's degree

of tallness is greater than mine, their degree of shortness will be less than mine, and so ensures that sentences like (1) will be true.

The hypothesis also explains the existence of cross polar anomaly. Since positive adjectives like *tall* are functions from individuals to positive degrees, negative adjectives like *short* are functions from individuals to negative degrees, and positive degrees are neither subsets nor supersets of negative degrees, the result of applying $\llbracket\text{-er}\rrbracket$ to a positive degree and a negative degree will be undefined. On the assumption that degrees on one scale are neither subsets nor supersets of degrees on another scale, cross modal anomaly can be given an exactly similar explanation.

2 Metasemantics for Formal Semantics

No doubt there is much more to be said. But it should already be plausible that Kennedy's view has the resources to explain the phenomena. But how does the explanation work? We know that the explanation involves assigning English expressions semantic values. But what does it mean to say that an expression has a certain semantic value?

I propose that we approach this question by way of a related *foundational* or *metasemantic* question. We have assumed that the semantic values of complex expressions are determined by the semantic values of their constituents and their structure. But how do the atomic expressions get their semantic values? This question is particularly pressing because it may seem that the sort of semantics we have been giving precludes a satisfactory metasemantics. How could a word like "tall" come to have a function as its semantic value – especially a function that maps objects to set-theoretic constructions far removed from ordinary thought? The usual philosophical stories rely some combination of the psychological states of speakers – for example, their intentions – and causal, evolutionary/teleological, or social factors. But none of these seems appropriate. Set theorists (and formal semanticists) may have relevant intentions towards the kinds of functions and sets that we have said are the semantic values of gradable adjectives, but two-year-olds – who readily master words like *big* and *tall* – do not. We do not causally interact with such sets and functions; and it is hard to believe that our cognitive systems are somehow designed by evolution to latch onto them.

So no traditional metasemantic story looks adequate for the kind of semantics that we are considering. But of course a good account of the nature of semantic theorising must be able to give a plausible metasemantic story. This is the first desideratum that our account of the nature of semantic theory must meet:

Metasemantic Desideratum An account of the nature of semantic theorising must make possible an account of how expressions come to have as their semantic values sets, functions, and similar abstracta.

A further observation may seem to exacerbate the problem. I have not specified precisely which sets and functions we are talking about. (I wrote, “a scale is understood to be a linearly ordered set of points, and a degree is an interval on the scale,” but of course this does not pick out a unique set.) The explanation can succeed in spite of this because any set with certain characteristics could do the job: all that we have assumed so far is the existence of a linear ordering, enough points on the scale that individuals that differ in height can be mapped to different degrees, and that the scale associated with *tall* and *short* is distinct from the scale associated with other adjectives (that generate cross-modal anomalies with *tall* and *short*).

One could make analogous points about the proposed explanation in a number of other ways. We have supposed that the semantic values of both gradable adjectives and comparative morphemes are functions. Functions are typically defined as sets of ordered pairs, and ordered pairs are in turn typically defined as sets of a certain kind. But there are multiple, equally adequate definitions: to take a trivial example, $\langle x, y \rangle$ can be defined as $\{\{x\}, \{x, y\}\}$, or as $\{\{y\}, \{x, y\}\}$. (Alternatively, the notion of an ordered pair, or indeed, the notion of a function, could be taken as a primitive.) These choices will make no difference to our explanations in semantics, but different choices would result in different entities being assigned as the semantic values of our expressions.

It is easy to see many places where a similar proliferation of candidate semantic values can arise. An account of the nature of semantic theory must be able to explain this proliferation:

Proliferation Desideratum An account of the nature of semantic theorising must explain (or explain away) the apparent fact that that there are many systems of abstracta that could do the same explanatory work.

One *prima facie* reason to see a problem here is that a seemingly similar proliferation of candidate set-theoretic definitions is widely thought to be problematic in another context. As Paul Benacerraf (1983) pointed out with respect to the project of reducing arithmetic to set theory, if one set-theoretic entity can play the role of (say) the number two, many others could do the job equally well. If numbers are set-theoretic objects, we ought to be able so say which set-theoretic objects they are. Is the number two $\{\{\emptyset\}\}$, or is it $\{\emptyset, \{\emptyset\}\}$? There is just no good answer to this question. Benacerraf

concluded that the attempt to identify numbers with sets is a sort of category mistake, and that sentences like “ $2 = \{\{\emptyset\}\}$ ” are either meaningless or false.

Why is the proliferation of candidates for reducing numbers to sets problematic? A first clue is that the project in this case is to say what numbers *are* – that is, what they are identical to. Since $\{\{\emptyset\}\}$ is not identical to $\{\emptyset, \{\emptyset\}\}$, at most one of them can be identical to the number 2. So the candidates are in competition: if we choose one, we are committed to saying that the others are wrong.

But there are other cases that do not have this feature, and in which the proliferation of candidates does not seem to be a problem. Suppose I measure the length of an object and discover it to be 6 inches; you measure the same object and discover it to be 15.24 centimetres. Suppose someone were to ask: which is the *real* length, 6 or 15.24? (Or, worse, which number is identical to the length?) That is a confused question. Lengths are not identical to numbers. Instead, we have set up a way of associating numbers and lengths. There are many such correspondences: feet-and-inches, centimeters-and-meters, etc. It may turn out that one or the other of these is the most useful for some particular purpose: for example, the metric system makes various kinds of calculation easier. But of course this does not entail that 15.24 is the “real” length (whatever that might mean). Even if I use meters, I need not say that someone who uses feet is wrong. (They may be impractical, unscientific, backward, American – but their claims are not false.)

Semantics is said to be the study of meaning. If we took a formal semantic theory to be making a claim about what meanings are (identical to) – for example, that the meaning of *tall* is a function from individuals to sets – then the proliferation of candidates would be a problem and formal semantics would be in trouble. But that should not be our attitude. Giving a formal semantic theory is much more like setting up a way of associating numbers with some physical quantity for the purposes of measurement.¹⁰

If giving a formal semantic theory is like setting up a system of measurement, the Proliferation Desideratum can be met. There exist many correspondences between numbers and objects, and many correspondences between set-theoretic entities and natural language expressions; arguably, these should be thought of as mathematical objects, existing independently

¹⁰There are a number of historical precedents for this idea. A number of philosophers – notably Stalnaker (1984) and Matthews (2007) – have developed the idea that assigning propositions to psychological states is like measurement. And the idea that a semantic theory is a system of measurement can be found in Davidson; e.g., Davidson (2001a, p. 130-3). (See Scharp (2013, ch. 7) for discussion.)

of human thought and activity. Setting up a system of measurement involves setting up a pattern of activity that exploits some such correspondence – for example, by describing procedures for assigning numbers to objects. Call such procedures a *system of measurement*. Since many systems of measurement will suit our purposes equally well, it is no surprise that there are many candidate units; and it is no surprise that there are many candidate semantic values, if semantic values are relevantly analogous to units of measurement.

The rest of this paper develops the analogy between semantics and measurement. Of course, semantics differs from paradigmatic cases of measurement, such as the measurement of length using a ruler, or the measurement of temperature using a thermometer. Measurement is associated with notions of quantity and scale; it typically involves assigning *numbers* to objects using *instruments*, while semantics does not involve quantity or scale in any normal sense, is not numerical, and does not (typically) use instruments. Three features of measurement are important for the analogy I want to press:

1. Measurement involves exploiting a correspondence between a phenomenon and certain abstracta (numbers or set-theoretic entities), so that features of the abstracta represent features of the phenomenon.
2. A system of measurement is prototypically set up by making a stipulation about some particular individual or individuals (e.g. stipulating that “one meter” is to pick out the length of a certain rod), and describing procedures by which the system can be extended beyond this stipulation.
3. Systems of measurement are set up and developed across particular periods in history, by agents whose knowledge of the domain to be measured may be very incomplete.

Whether or not to reserve the word “measurement” for cases involving numbers and instruments is a relatively uninteresting question of terminology; my view could be developed in other terms (for example, in terms of Swoyer’s (1991) notion of structural representation) as long as these three points of analogy are attended to.

It is easy to imagine the task of the semanticist (or of the scientist generally) as mostly descriptive: the facts are out there, and our task is to state or describe them. But we are already in a position to see that the situation is not so simple. Prior to investigation, we may not have adequate means for representing the relevant facts. In this case, the role of the theorist will go beyond merely stating and describing: she must *set up* a system of represen-

tation. I will go on to discuss how this can be done. It will clarify the issue if we first return to the Metasemantic Desideratum.

3 Metasemantics, Targets, and Perspectives

If assigning semantic values is relevantly like setting up a system of measurement, then we can partially meet the Metasemantic Desideratum: it is in virtue of the choices of theorists that expressions have the semantic values they do. But this answer is only partial. The choices were not mere whim. Why have semanticists made the choices they have? Slightly more precisely: we have portrayed the semanticist as describing a certain set-theoretic structure and claiming that it can be used to represent certain features of language. But exactly what is it being used to represent? (If a semantic theory is giving something like units for measuring, what are we measuring?) Call this *the target question*.

There is a good deal of disagreement about the right answer to the target question, both among philosophers and among linguists. (Is semantics in the business of characterising a mental faculty or organ (Chomsky, 1995)? Or characterising social facts, such as conventions (Lewis, 1983)? Or of describing the information is encoded by sentences in contexts (Soames, 1989)? Or of giving a theory that would put someone who knew it in a position to interpret utterances (Davidson, 2001b)? And no doubt there are other possibilities.) Obviously, we cannot resolve these debates here, and even if we had an answer, the fact that progress in semantics is possible despite disagreement about the nature of the target is itself remarkable and in need of explanation. How can we make progress in theorising about a phenomenon, when we have radically different views about the phenomenon we are characterising (so that at least some of us have radically false views about the nature of the target, and even those with views in the right ballpark may exhibit significant errors or confusions)? An account of semantic theorising should be able to answer this question:

Target Ignorance Desideratum An account of the nature of semantic theorising should make it possible to understand how we have made progress in semantic theorising despite radical disagreement and ignorance on the target of that theorising.

It may seem surprising that we have made progress in semantics despite ignorance and disagreement, but there are precedents in other sciences. Consider the scientist investigating temperature in the eighteenth century. The target phenomenon was not at all well understood; for example, on

some views cold was not merely the absence of heat, but an independent force; others failed to distinguish temperature and heat. There were deep disagreements between proponents of various versions of the caloric theory and their opponents. There was no scale for measuring temperature, and it was not even clear that such a scale could be developed (a point we will return to below). As a consequence, only a few kinds of facts about temperature that could be represented (notably, relational facts (for example, about what is warmer than what)).

We can contrast the position of the eighteenth-century theorist of temperature with the position of a theorist trying to understand our current systems for measuring length. In this case, the target is relatively well understood, and there is little relevant disagreement about its nature. We have well-established systems of measuring length that allow us to represent a wide range of facts about lengths. The question in this case is not: how can we set up a system of measurement? It is rather: how does our extant system of measurement work?

When we think about the representation of a given phenomenon, we begin with background knowledge and abilities, which may be more or less extensive. Typically, we will fall somewhere on an epistemic spectrum, the ends of which – what I call *the Perspective of Innocence* and *the Perspective of Experiences* – are roughly illustrated by the cases of temperature in 1700 and length today.

Perspective of Innocence	Perspective of Experience
Target not well understood (theories of target may be rudimentary, confused, and/or false; target may not be clearly distinguished from related phenomena)	Target relatively well understood
Disagreement among experts about the nature of the target	Little relevant disagreement among experts about target
No extant scale of measurement, or clear means of establishing one	Already extant scale of measurement
We can make precise and accurate judgments about only a narrow range of facts about the target (because we lack means of representing other sorts of facts, relevant instruments, know-how, etc.)	We can make precise and accurate judgments about a wide range of facts about the target

In my view, though our investigations in semantics have made very great progress, we are in crucial respects still close to the Perspective of Innocence.

We have already pointed out that there is little agreement about the nature of the target of semantic theorising. And even among proponents of a specific view of the target, it is often agreed that investigation is at a relatively early stage (e.g., Yalcin (2014, p. 49)).

From the Perspective of Innocence, our question is a practical one: how can we set up a system of measurement? But in order to answer this question, it will be useful to look at the kind of understanding of a measurement that we can get from the Perspective of Experience. I turn to this task in the next section.

4 Measurement from the Perspective of Experience

Suppose that we are interested in understanding a system of measurement for length. We already know a good deal about length, and we already have good scales for measuring lengths. Why do these scales work?

The mathematical discipline of measurement theory is the natural place to look for an answer. The measurement theorist offers to prove that for any set of objects, operation \oplus on those objects, and relation R on the set, if certain constraints are met (we can ignore the details), then there is a way of assigning numbers to objects such that for any objects a (to which we assign a number n_a) and b (to which we assign n_b), a stands in R to b iff $n_a \geq n_b$, and $a \oplus b$ is mapped to $n_a + n_b$; this is sometimes called a *representation theorem*. Moreover, the measurement theorist offers to prove that any two such mappings differ only by a multiplicative factor (i.e., for any mappings A and B , there is a number n such that for any object x , A maps x to a iff B maps x to na); this is sometimes called a *uniqueness theorem*.

How does this bear on measuring length? The idea is simple: consider the set of physical objects, the relation of being at least as long as, and the operation of putting two objects immediately next to each other end-to-end. Then, by the measurement theorist's result, we can assign numbers to objects in such a way that if a is at least as long as b , then the number assigned to a is greater than or equal to the number assigned to b , and the combination of a and b placed next to each other end-to-end is mapped to the sum of the number assigned to a and the number assigned to b . All we need to generate such a mapping is a unit; in other words, we need an object (like the "standard meter" bar) that we can assign the number 1. Plausibly, we can regard numbers considered under such a mapping as representing facts about length.

The uniqueness result explains why our choice of unit is arbitrary (practical considerations aside). Whether we choose feet or meters, the measure-

ment theorist's result will go through; which we choose is a matter of convenience. But we have made another arbitrary choice, which is so natural that it may slip by unnoticed: we have chosen to let the length of two objects concatenated be the sum of their individual lengths; i.e., we have chosen map $a \oplus b$ to $n_a + n_b$. But this is not mandatory. We can (with a suitable modification to the theorem to be proved) associate our operation on the domain of entities measured with a different operation on numbers (see Ellis (1968, pp. 78-81)). Suppose we introduce a new kind of unit for length – *squards* – with the stipulation that for any object x , x 's length in squards is the square of its length in yards. Then we could not associate the concatenation operation with addition: instead, we would need want our mapping to obey the following constraint: if a is mapped to m_a and b to m_b , then $a \oplus b$ is mapped to $(\sqrt{m_a} + \sqrt{m_b})^2$. This is somewhat unwieldy for many ordinary purposes, but may be useful for others; and in any case, is adequate to capture all of the facts about length that can be captured using our ordinary units.

As the example of squards shows, there are indefinitely many ways of mapping objects to numbers, even once we have decided that the relation of being greater than or equal to corresponds to the relation of being at least as long as, and picked a unit length. We have singled out a particular mapping by imposing an additional constraint: in the case of metres or feet, we have stipulated that the combined length in metres or feet of two objects placed end-to-end is the sum of their individual lengths (in meters or feet); in the case of squards, we have stipulated that the combined length of two objects meets the more complicated constraint just described. Suppose we have stipulated (e.g.) that “one meter” is to name the length of a certain bar. The idea that we need to impose an additional constraint on our mapping in order to extend the scale beyond this stipulation will be important in what follows; we will refer to this as a *scale constraint*.

So we can use the measurement theorists' result to understand how scales of measurement of length represent lengths. But this requires substantial knowledge. First of all, we need to know what set of objects, relation, and operation to consider. As a first step, we need to know what kind of objects have lengths (physical objects and not, say, numbers). We also need to understand how to put two objects end-to-end so as to evaluate their combined length. Even in the case of objects that are relatively simple to handle – rigid rods, say – we need to know to put them so that only their ends are touching, rather than (say) so that they overlap. Gaining even a basic grasp of how to think about this operation across the full range of physical phenomena that have lengths would require a huge amount of knowledge (probably including knowledge, such as know-how, that is difficult to articu-

late explicitly).

We can apply the resources of measurement theory only when we already know something of the domain we are measuring. In the case of length, this comes easy: probably most if not all adults have most of the knowledge described, at least implicitly. But we have not always had scales for measuring length. Suppose we lacked such a scale; suppose we were in this respect closer to the Perspective of Innocence. How could we create such a scale for the first time?

4.1 Introducing a Scale of Measurement

An obvious first step is to pick an object or phenomenon to serve as a standard unit: if we are unsophisticated, it may be the king's foot; if we are more sophisticated, the length of a certain platinum-iridium bar, or the distance travelled by light in a vacuum in a precisely specified amount of time. To a first approximation, we use the chosen object or phenomenon to fix our scale by stipulating that (e.g.) "one meter" is to name its length.¹¹

It may seem that our choice of a standard unit is entirely arbitrary, and it is certainly true that we have many options. But some choices are better than others: a snake, or a rubber band, would serve badly as a standard (if at all). What exactly is the problem with using an object that is highly variable in length as a standard? Consider, for example, the stipulation that "one meter" is to name the length of a particular rubber band. There are several ways such a stipulation might be intended, and precise nature of the problem will depend on exactly how we interpret it. If what is intended is that an object is one meter long at t just in case it is the same length as the rubber band at t , then the problem is practical: the length of a given object in meters will fluctuate as the rubber band is stretched, making the determination of lengths, the statement of scientific laws, calculations, etc. very difficult. If we stipulate that an object is one meter long at t just in case it is the same length as this rubber band is now, then we will need some way to determine whether some future object is the same length as the band is now (and the obvious way to do this involves finding a less elastic object that is the same length as the band, which would then serve as a *de facto* standard).

¹¹What kind of thing is a length? Realists maintain that quantities like length exist independently of our systems and practices of measurement; one kind of anti-realist argues that our measurement operations are partially constitutive of length. The stipulation we have described seems to presuppose a realist view, since it seems to presuppose that the lengths of objects are out there to be named. In fact, the situation is probably not so simple, and I think that the discussion to follow could be reconstructed in an anti-realist way. But for the sake of simplicity, I will not resist the temptation to speak with the realist.

We who are familiar with the stretchiness of a rubber band would most likely not try to use it to introduce a unit of length – or if we did, we would have a sophisticated interpretation of the stipulation (such as those considered in the previous paragraph) in mind. But consider the person who stipulates that “one foot” is to name the length of the king’s foot, either unaware of or simply ignoring the fact that the king’s foot swells a bit at the end of a long day of throne-sitting. Such a person need have no sophisticated interpretation of the stipulation in mind; they are simply presupposing that the length is constant. And we can imagine a suitably ignorant or deluded stipulator, who is unaware that rubber bands stretch, doing the same.

When the stipulation “Let ‘one meter’ name the length of this rubber band” is made in these circumstances – i.e., made by someone who is presupposing that the length of the rubber band is constant – it is plausible that the stipulation simply fails. The stipulator has not done enough to pin down the meaning of “one meter”. Absent some further stipulation (or something else – perhaps a pattern of use – that does similar work), sentences involving uses of “one meter” will lack truth values.

Too much variation in length is problematic, but our stipulation need not involve an object that is absolutely constant. The stipulation that “one foot” is to name the length of the king’s foot can be used to introduce a unit for measuring length; its variation in length is small enough to be insignificant for many practical purposes. It seems perfectly reasonable to measure medium-sized dry goods in feet following such a stipulation (though the foot would be a problematic unit in cases where a great deal of precision was called for), and claims like “I am six feet tall” might be useful and (plausibly) even true. The fact that we can rely on standards that are not entirely fixed or constant is critically important for most – perhaps even all – of our ordinary and scientific scales of measurement. The length of a carefully-made platinum-iridium bar can vary slightly, though these variations matter only for extremely precise scientific work; the speed of light may be constant, but using it to define a standard of length means relying on a unit of time, the constancy of which may be questioned (Tal, 2011; Scharp, ms).

Many successful stipulations take place against a background of ignorance or ignoring of the fact that the phenomenon we are using in our stipulation is not truly constant with respect to the feature we want to measure. It is natural to regard this as a kind of idealisation: we are ignoring variations in length, or representing the king’s foot as being constant in length when it is not.¹² Suppose that we stipulate that “one foot” is to name the length

¹²Some authors distinguish between idealisation (in which an object is represented as hav-

of the king's foot, or that "one meter" is to name the length of a certain bar. Suppose further that the stipulation succeeds: we go on to use feet or meters to describe the lengths of various objects, and that the variation in length of the king's foot or the meter bar causes no practical problems for our applications. How, then, should we view claims like "This tree is three feet tall"? There are a number of options in the literature on related problems (such as vagueness and the problem of the many). Perhaps it is metaphysically indeterminate exactly which length is a foot. Perhaps subtle features of our linguistic practice (perhaps including those that go beyond stipulation) make it the case that "one foot" picks out some precise length. Perhaps such claims are strictly speaking false, though useful for some reason.

One familiar account – supervaluationism – provides an appealing model. (I take no stand on whether it is more than a useful model.) Suppose that the king's foot is exactly 0.30 meters at the beginning of the day and exactly 0.31 meters at the end of the day. The supervaluationist points to many possible candidates for being the precise length picked out by "one foot": it might be 0.300 meters, or 0.301 meters, or 0.302 meters, and so on up to 0.31 meters. We can evaluate sentences on the supposition that one or the other of these candidates is correct. For example, on the supposition that "one foot" picks out 0.300 meters, and that a given tree is 0.903 meters tall, the sentence "This tree is exactly 3 feet tall" is false; on the supposition that "one foot" picks out 0.301 meters, that same sentence is true. The supervaluationist then claims that a sentence is true simpliciter just in case it is true under every relevant supposition, false simpliciter just in case it is false under every relevant supposition, and neither true nor false otherwise.

The model is useful for our purposes because it predicts that a unit introduced by a stipulation of the type described involving an object that is more variable in length will be useful in fewer applications than a unit introduced by a stipulation involving an object that is relatively less variable in length. For example, a stipulation involving a rubber band will be useful for very little; a stipulation involving the king's foot will be useful for ordinary measurement but not for scientific work; a stipulation involving a platinum-

ing a feature it in fact lacks, or as lacking a feature it in fact has) and abstraction (in which a representation is silent on whether an object has a certain feature). I do not want to insist that what goes on in such stipulations is a matter of idealisation in this sense, rather than abstraction. But I would want to insist that if we describe such stipulations as involving abstraction, it is abstraction on a matter of central relevance to the stipulation: we are silent on whether the king's foot varies in length, even though whether the stipulation succeeds in introducing a unit of measurement that is adequate for some purpose will depend on the extent to which the king's foot in fact varies in length.

iridium bar will be useful for ordinary and many scientific purposes; etc. Roughly, different purposes require different standards of precision. In an ordinary context, we will regard an utterance of “The tree is 3 feet tall” as true iff the tree is close enough to three feet tall – in other words, as roughly equivalent to an utterance of “The tree is between 2’9” and 3’3” tall”. If we have introduced “feet” by a stipulation involving the king’s foot, such an utterance could easily come out true. (Suppose that the tree is 0.915 meters tall. Then on the supposition that “one foot” picks out 0.300 meters, the tree will be 3.05 feet – less than 3’1”. On the supposition that “one foot” picks out 0.310 meters, the tree will be 2.95 feet – more than 2’11”. So on every relevant supposition, “The tree is between 2’9” and 3’3”” will come out true; so the supervaluationist predicts that utterance is true simpliciter.) In a scientific context, we may regard an utterance of “The tree is 3 feet tall” as being roughly equivalent to something like, “The tree is within a tenth of an inch of 3 feet tall”. If we have introduced “one foot” using the king’s foot, this sentence will be neither true nor false. (Suppose that the tree is 0.915 meters tall. Then on the supposition that “one foot” picks out 0.300 meters, the tree will be 3.05 feet – more than a tenth of an inch over 3 feet – so on this supposition, the sentence is false. But on the supposition that “one foot” picks out 0.305 meters, the tree will be 3 feet exactly, and the sentence is true. Since on some suppositions the sentence is false and on others it is true, the supervaluationist predicts that it is neither true nor false simpliciter.) But if we have introduced “one foot” using a carefully-made platinum-iridium bar the sentence could easily come out true simpliciter.

In many cases, the history of measurement has been a history of working toward ever more precise scales. And we can see why: a scale introduced by a stipulation involving the platinum-iridium bar works when a scale introduced by a stipulation involving a foot does not. But the king’s-foot scale works sometimes; and this is important, because in early stages of investigation – from the Perspective of Innocence – a scale of this sort may be all we have. It is to cases of this type that we now turn.

5 Measurement from the Perspective of Innocence

Finding objects whose length remains constant – or at least, constant enough for our purposes – is easy, in part because of our substantial background knowledge about length. But how is it possible to introduce a scale of measurement for a phenomena that we understand much less well? I propose to gesture toward an answer to this question by examining the case we used to introduce the Perspective of Innocence: temperature. (My discussion in the

following paragraphs is deeply indebted to the excellent Chang (2004).)

The first problem is finding objects constant in temperature that we can use to pin down the scale. As Chang points out, in the case of temperature it was not obvious a priori that there are any “fixed points” – “phenomena that could be used as thermometric benchmarks because they were known to take place always at the same temperature” (2004, p. 9). Through the seventeenth and eighteenth centuries, a number of proposals were mooted (including the temperatures of melting butter, the cellars of the Paris observatory, and the human body, the last of which was advocated by both Newton and Fahrenheit) (2004, p. 10). Even the reasonably reliable boiling and freezing points of water vary very considerably depending on factors such as the amount of air dissolved in the water, and the nature of the vessel in which the water is heated or cooled (as well as on atmospheric pressure); the still more reliable temperature of steam depends on the presence of dust in the air (2004, ch. 1).

We saw in the previous section that our stipulation of a unit for measuring length need not involve an object that is absolutely fixed in length – the stipulation is (or at least, can be) *idealised*. The same is true for the stipulation of fixed points in the scale of temperature. These idealisations can give us some insight into a puzzle about the introduction and development of scales of measurement.¹³ When we began to theorise about temperature, we were near the perspective of innocence. But we need significant knowledge to introduce a scale of measurement; and from the perspective of innocence, we do not have such knowledge. The puzzle is that we also seem to lack the means to gain it. How can we tell whether a given phenomenon is a fixed point? Well, we could tell if we had a good thermometer. But we can’t build one without a good theory – something else that we don’t have in the Perspective of Innocence. And we can’t even state a theory without a scale for measuring temperature, much less develop and test it without good instruments. So it looks like we have a sort of catch-22: we can’t develop our theories without a scale of measurement, but we can’t set out a scale of measurement unless we already have a theory.

Of course we did make progress: every drugstore now carries reliable thermometers calibrated to standard scales, and every elementary physics

¹³In the case of temperature, further difficult problems arise when we consider how to formulate a scale constraint to enable us to extend the scale beyond the fixed points (and this may involve further idealisation), but this is among the many complications that considerations of space force us to gloss over. For more thorough and sophisticated treatments of this sort of puzzle, broadly consistent with the spirit of the view developed here, see Chang (2004); van Fraassen (2008).

textbook contains sophisticated theoretical discussion of temperature and heat. How did we move from the Perspective of Innocence to the Perspective of Experience? We now have enough information to isolate two important factors:

1. Although there was much that we did not know about temperature, and what little knowledge we had was not theoretically articulated, it isn't as though we had no information or means of gaining it whatsoever. We can come to know some basic facts about what is warmer than what via our sense of touch, and we can observe the behaviour of various substances as they are heated.
2. We can introduce units of measurement via idealised stipulations even when we have not managed to pinpoint truly fixed points or a perfect thermometric substance.

These facts help explain how we were able to bootstrap our way from theoretical rags to riches. By making use of the knowledge described in (1), we were able to make a reasonable guess about possible, rough fixed points and thermometric substances. Using this knowledge, we can make idealised stipulations of the sort described in (2) to introduce a scale. Because of our idealisations, the scale may only be useful in circumstances where great precision is not needed. But it is a place to start: with it, we can begin to build instruments and more sophisticated theories, and these will in turn put us in a position to set up better scales of measurement. We do not build on a firm foundation; it would be more accurate to say that we are constantly revising an idealised one.

These facts put us in a position to meet the Target Ignorance Desideratum. It is possible to set up a scale of measurement despite lack of an agreed background theory because the setup of the scale can involve stipulation based on relatively theoretically neutral background knowledge, which is likely to be widely shared even among theorists who disagree on theoretical matters: all we need is a rough idea of what the fixed points are and what the scale constraint could be. With these in place, we can make idealised stipulations that can be used to formulate and test our theories, including our theories about the nature of the target.

Because the relevant stipulations are (at least relatively) theory neutral, they can be endorsed by theorists with different theoretical views. Of course, in some cases it might turn out that theorists with what seem to be radically different views are in fact theorising about different phenomena. (Theorists might have taken themselves to disagree about a common thermal phe-

nomenon when in fact some were trying to measure temperature and some trying to measure heat.) But not every case is like this: theorists can make progress on developing a scale for measuring a single phenomenon while disagreeing about the nature of that phenomenon. So while there is room to admit that some disputes about semantics involve theorists talking past each other (*perhaps* Chomskians are just interested in something different than Lewisians), there is also room to admit genuine disagreements and genuine progress on a common phenomenon despite them.

6 Semantics as Measurement: Explanation

In the previous sections, we noted some factors that are important in setting up a scale of measurement from the Perspective of Innocence. The rest of this paper shows how these factors play a role in semantic theorising. In order to get clearer on exactly how semantics is like setting up a scale of measurement, we need to return to the degree theoretic semantics for gradable adjectives discussed above, with an eye toward understanding the explanations of the entailment and anomaly facts that this semantics offers. On the view that setting out a semantic theory is like setting up a scale of measurement, these explanations may seem quite puzzling. How can setting up a scale of measurement be explanatory?

It is true that we often formulate explanatory claims in terms of units of measurement. Why was he scalded when he dropped the cup of tea? The tea was 100 degrees C! But this explanation relies on empirical background knowledge: that 100 degrees C is hot enough to scald. Simply setting out the Celsius scale for measuring temperature is not enough to generate the explanation. The semantic explanation seems different; the explanation of cross-polar anomaly (say) does seem to be generated in a more direct way by the semantics, so that merely setting out the semantic theory is enough to generate the explanation.

But the situation is not so puzzling when we reflect on the kind of stipulations that go into setting up a scale of measurement. Setting out a scale of measurement can be explanatory because of the information that is captured by these stipulations. The situation is perhaps clearest when we rely on a scale constraint, such as the association between concatenating objects and addition that we relied on in introducing the scale of length. Consider an example. Why is the length of an American football field 120 yards? Well, once we know that the main field of play is 100 yards long, and immediately adjoined to each end is an end zone that is 10 yards long, the explanation

is obvious,¹⁴ and flows immediately from the scale constraint: if an object of length n yards is adjoined to an object of length m yards, then their combined length is $n+m$ yards; and $100+10+10 = 120$. The explanation turns in part on the nature of length, but also in part on the unit constraint that we have chosen; if we were measuring in squards, our explanation would have to take a different form.

Setting up a scale for measuring length, or for measuring temperature, involved two kinds of stipulations:

1. Stipulations about “fixed points” – i.e., stipulations linking specific values to phenomena believed to be constant with respect to their length or temperature, such as “Let ‘one foot’ name the length of the king’s foot”, or “Let ‘o degrees’ name the temperature at which water freezes.”
2. Stipulation of a scale constraint, to expand the scale beyond the fixed points.

How do these two kinds of stipulation match up to what it takes to set out a semantic theory? Consider again the explanation of cross-polar anomaly discussed above. We begin with the observation that sentences like “Derek is taller than Brian is short” are anomalous: they are impossible to interpret, “sound funny”, etc. We set ourselves the task of assigning semantic values to “Derek”, “tall”, “-er”, “Brian”, and “short”, which are such that given the syntactic structure of the sentence and standard semantic composition rules, the sentence will not be assigned a semantic value. When we can show that it follows from the syntactic structure of the sentence, and our stipulated semantic values and composition rules that the sentence has no semantic value – for example, because we can show that (by the composition rules and the syntactic structure) $\llbracket \text{Derek is taller than Brian is short} \rrbracket = \llbracket \text{-er} \rrbracket (\llbracket \text{Derek is tall} \rrbracket) (\llbracket \text{Brian is short} \rrbracket)$, that (by deriving $\llbracket \text{Derek is tall} \rrbracket$ and $\llbracket \text{Brian is short} \rrbracket$ from their syntactic structure and the semantic values of their atomic parts via the composition rules) $\llbracket \text{Derek is tall} \rrbracket$ is neither a subset nor a superset of $\llbracket \text{Brian is short} \rrbracket$, and that therefore (by the semantic value of $\llbracket \text{-er} \rrbracket$), $\llbracket \text{-er} \rrbracket (\llbracket \text{Derek is tall} \rrbracket) (\llbracket \text{Brian is short} \rrbracket)$ is undefined – then we regard the phenomenon as explained.¹⁵

¹⁴Of course this is not to say that there are no other kinds of explanation – e.g., historical – of the length of an American football field. But it seems clear that what follows is one type of explanation.

¹⁵As before, this is not to say that no other kind of explanation is possible. (We have already asked for a metasemantic explanation (of how the atomic components get their semantic values).)

If we take the syntactic structure as given, the explanation has at least three moving parts:

- (A) The assignment of semantic values to atomic expressions. (In particular, we assigned people as the semantic values of names, and functions of various kinds as the semantic values of *tall*, *more*, etc.)
- (B) The description of composition rules that determine the semantic of complex expressions given the semantic values of their parts and their structure. We have relied on standard assumptions here, without making them entirely explicit: the idea is that the semantic value of a complex expression is derived by function application from the semantic values of its constituents, and that the relevant notions of complexity and constituency are syntactic.
- (C) The association of semantic values (or the lack of a semantic value) with various phenomena to be explained. Again, we have relied on standard ideas here, and so left this mostly implicit; but the kind of examples I have in mind are the claim that a sentence is anomalous if it has no semantic value, or that one sentence entails another if the semantic value of the latter is 1 when the semantic value of the former is 1.

Now we are looking for something that plays a role like the stipulation of fixed points, and something that plays a role like the stipulation of a scale constraint (that enables us to extend the scale beyond the fixed points). Clearly (B) is not a stipulation of a fixed point. It does not fix the semantic value of any expression; instead, it describes how the semantic values of some expressions are related to the semantic values of other expressions. In other words, it constrains the possible assignments of semantic values: if we have semantic values for atomic expressions, then (B) determines the semantic values of complex expressions (given their syntactic structure); if we have the semantic values for complex expressions (and their syntactic structure), then (B) puts limitations on the possible semantic values of atomic expressions. So our composition rules help us fix the semantic values of some expressions on the basis of the semantic values of other expressions. This sounds very much like the job of a scale constraint; and this is exactly the role I take composition rules to play.

A scale constraint puts us in a position to expand a scale beyond the stipulation of particular fixed points. What, then plays the role of the stipulation of fixed points? Consider the explanation of cross-polar anomaly as I described it. There we began with a goal – assigning sentences like

“Derek is taller than Brian is short” no semantic value – and we sought to assign semantic values to atomic expressions that generated this result. This strongly suggests that the association of anomalousness with having no semantic value is acting as a fixed point, and that the role of the composition rules is to constrain our assignment of semantic values beyond this.

So the role of one of the stipulations described in (C) is to set a fixed point. But (C) also describes another stipulation – that one sentence entails another just in case the semantic value of the latter is 1 when the semantic value of the former is 1. Like the stipulation described in (B), this is not plausibly thought of as the stipulation of a fixed point, because it does not fix the semantic value of any expression. Instead, it constrains the overall pattern of assignments: if we assign 1 to “Derek is taller than Brian”, then we must assign 1 to “Brian is shorter than Derek”.

What, then, of the assignment of semantic values to atomic expressions described in (A)? Taking on board what we have so far, our goal is to make an assignment of semantic values that (in combination with the composition rules) entails that sentences get assigned: i.e., that anomalous sentences have no semantic value, and that sentences get assigned 1 only if the sentences they entail get assigned 1. Plausibly, there are many ways of doing this. (We have already mentioned that there are many sets that could serve as scales for use in the semantics of gradable adjectives.) Crucially, however, our choices are *mutually constraining*. Once we have stipulated that **[[tall]]** is a function that maps individuals onto degrees on a certain scale, we must also stipulate that **[[short]]** maps individuals onto degrees on that same scale, and we must stipulate that **[[beautiful]]** maps individuals onto degrees of a different scale. Or, to take another kind of example, one way to deliver the result that “Derek is taller than Brian is short” is assigned no semantic value would be to stipulate that “Derek” is assigned no semantic value; but we are prevented from taking this easy way out by the fact that we need other sentences involving “Derek” (such as “Derek is tall”) to be assigned a semantic value.

So the semantic values of atomic expressions play something of a mixed role. In the end, we do want the semantic values of atomic expressions to be fixed points, since we want atomic expressions to contribute the same thing to every complex expression in which they occur. But our assignment of semantic values to atomic expressions is not just arbitrary, in the way that assigning 100 (or 232, as the case may be) to the temperature of boiling water is arbitrary. We are constrained in each choice we make by the choices we have made elsewhere. Mapping out these constraints systematically would be an interesting task, but is not one we can undertake here: it still remains

to revisit the problem of radical contextualism with which we began.

7 Idealisation and Metasemantics

Recall that the radical contextualist points out that whether an utterance of a sentence like “The leaves are green” seems true or false depends on abstruse aspects of the context in which it occurs, aspects so varied that it is implausible to suppose that they could be formalised in the kind of semantic theory we are considering. How exactly does this observation bear on view of semantics developed in this paper?

To simplify discussion, let’s assume that the semantic values of sentences are truth values. (As we saw early in this paper, it is plausible that some version of the objection can be directed even against approaches to semantics on which this is not so.) Then the radical contextualist’s objection is naturally construed in something like this way: truth values of sentences vary depending on context in ways that are hard to systematise; so they are not good candidate fixed points. How should the formal semanticist respond?

A first response would be to deny that the semantic values of sentences are in general meant to be fixed points. In the explanations of cross-modal and cross-polar anomaly, the only fixed point related to the semantic values of sentences was the stipulation that anomalous sentences have no semantic value. And it seems plausible that anomaly is less susceptible to contextual manipulation than truth and falsehood. Set aside the (relatively easy to isolate (see footnote 9, above)) readings of the relevant sentences on which they do not seem anomalous. It seems hard to imagine contexts in which the sentences (on their ordinary readings) sound okay.

This is suggestive, but hardly constitutes decisive evidence that anomaly is a context-insensitive fixed point. But even if it were, there would still be a problem, because (as we have seen) there are fixed points elsewhere in the system. In particular, it is plausible that the assignment of semantic values to atomic expressions is a fixed point (albeit a fixed point assigned under constraint). And the radical contextualist should take her arguments to tell against this assignment as well. The crucial fact here is that (except perhaps in certain special cases, such as the treatment of quantifier expressions), that the assignment of semantic values to expressions uses the object language vocabulary that the radical contextualist is claiming to be context sensitive in the metalanguage; in particular, one typically assigns as the semantic value of an expression a set-theoretic entity that is specified *using that expression*. For example, an introductory textbook might begin with semantic values like these:

- (9) $\llbracket \mathbf{runs} \rrbracket = \lambda x.x \text{ runs}$
 (10) $\llbracket \mathbf{green} \rrbracket = \lambda x.x \text{ is green}$

Why is the use of context-sensitive vocabulary in the metalanguage problematic? Consider a less contentious example. Suppose I propose the following lexical entry for “that”:

- (11) $\llbracket \mathbf{that} \rrbracket = \text{that}$

Now we face a dilemma. Either the context in which we set out our semantic theory contains a enough shared beliefs, intentions, observable features, and so on, to fix the meaning of “that” as used on the right-hand side of the identity (as might be the case if I point at a certain salient object – say, my desk lamp – with appropriate Gricean intentions, and you are in a position to recognise this); or it does not (as might be the case if I do not point, have no relevant intentions, and no object is suitably salient). The problem with the first horn is (roughly) that our theory will falsely predict that every use of “that” picks out my desk lamp, and will make correspondingly wacky predictions about sentences involving “that”. This seems like a serious problem; so let’s turn our attention to the second horn.

The problem with the second horn is that if context does not fix any content for “that” as used on the right-hand side of the identity, then (11) fails to express anything; it cannot be used to make a stipulation at all. The version of the radical contextualist objection that we are considering has it that attempted stipulations like (10) fail in exactly this way. Ordinary conversations take place in a shared environment, against a rich background of shared beliefs, desires, intentions, and so on; and this complicated constellation of attitudes and environmental factors (or at least certain parts of it) plays an important role in resolving context sensitivity. For example, in the case of the green leaves, it is plausible that some combination of Pia’s interlocutors’ purposes, her knowledge of these purposes, the fact that some leaves but not others would satisfy these purposes (for example, because painting leaves changes their appearance but not their biochemistry) and her knowledge of this fact, her intentions, her interlocutors’ knowledge of these intentions, and so on play a role in determining whether or not painted leaves count as “green”. Without a background like this, a use of “green” would have no determinate content at all. But (the objection continues) the context in which we introduce a semantic theory is highly unusual. We are unlikely to have any particular intentions as regards whether (say) painted leaves count as “green”; and more generally, there is nothing in such a context that could make it the case either that painted leaves count as “green” or that they fail

to count as “green”. So in attempting to make a stipulation like (10), we are failing to pick out a function.

The objection has it that a stipulation like (10) is much like a stipulation of (12) (made against a presupposition that the length of the band constant):

(12) one meter = the length of this rubber band

The problem is that since the length of the rubber band varies, the expressions “the length of this rubber band” seems to fail to pick out a particular length (as long as we are presupposing that the length is constant, rather than trying to make one of the sophisticated stipulations described in section 4.1 above).

We can make a similar argument about a stipulation like (13):

(13) one foot = the length of the king’s foot

Despite this, the “foot” stipulation can be used to introduce a scale of measurement suitable for many everyday uses. On the view described above, this is because we treat (13) as an idealised stipulation: we idealise away from the foot’s variation in length.

My view is that we should view a stipulation like (10) as idealising away from contextual variation in much the same way. In the case of the king’s foot, we modeled this idealisation supervaluationally. Now we can make a precisely analogous move with respect to semantic value and a stipulation like (10). If the context in which we set out our semantic theory does not determine a single meaning for “green”, then no single function is described by “ $\lambda x.x$ is green”; there are various candidate functions that might be picked out by that expression, compatible with whatever metasemantic features obtain in the context. So we can supervaluate over these candidates. A sentence like (14) will come out true just in case every candidate function maps Kermit to 1, false just in case no candidate function does, and truth-valueless otherwise:

(14) $[\lambda x.x \text{ is green}](\text{Kermit}) = 1$

This is a promising start, but we are not done. A stipulation like (10) can introduce a scale that is useful for many purposes, but a stipulation like (12) cannot. In our supervaluational model, this is explained by the fact that if “meter” is introduced by a stipulation like (12), few ordinary sentences involving “meter” will be true or false, and many will be truth-valueless. The radical contextualist should object that the contextual variation is so great that most or all claims like (14) will be ruled truth-valueless. (For example,

suppose that our context of stipulation leaves it open whether things that are green on the inside, or things that are green on the outside, count as “green”. Assuming that Kermit is an ordinary frog, green only on the outside, then (14) will be truthvalueless.) Then (the objection continues) a semantic theory introduced by stipulations like (10) will be relevantly like a system for measuring lengths introduced by (12) – useless.

Scales introduced by idealised stipulations are useful for some purposes but not for others. Whether the present objection succeeds in showing that semantic theory is useless depends very much on what we propose to use it for. Consider the explanations produced by the degree-theoretic semantics for gradable adjectives. For example, our explanation of cross-polar and cross-modal anomalies turned on the claim that sentences like “Derek is taller than John is beautiful” will not be assigned semantic values, and this claim can be derived entirely from claims that are true on every candidate interpretation, hence true simpliciter on the supervaluational theory we are considering. Assume that our syntactic structure and composition rules deliver something like (15):

$$(15) \quad \llbracket \text{Derek is taller than John is beautiful} \rrbracket = \llbracket \text{-er} \rrbracket (\llbracket \text{tall} \rrbracket (\llbracket \text{Derek} \rrbracket)) (\llbracket \text{beautiful} \rrbracket (\llbracket \text{John} \rrbracket))$$

By the relevant lexical entries, we know that $\llbracket \text{tall} \rrbracket (\llbracket \text{Derek} \rrbracket)$ will be assigned a degree on the scale of height, and $\llbracket \text{beautiful} \rrbracket (\llbracket \text{John} \rrbracket)$ will be assigned a degree on the scale of beauty. Now given the lexical entry for “-er” described in (7), all we need to show that $\llbracket \text{Derek is taller than John is beautiful} \rrbracket$ is undefined is the claim that no degree on the scale of height overlaps a degree on the scale of beauty. But our choice of entities to act as scales is arbitrary – essentially a matter for stipulation. So we can be sure – we can stipulate – that the scales will not overlap, regardless of how the context-sensitive expressions “height” and “beauty” that we use in the metalanguage are interpreted. Facts about entailment are similar; the claims needed to derive the entailment facts are true on any interpretation of the context-sensitive (because we have built them in to every interpretation by stipulation), so true simpliciter.¹⁶

So the first conclusion I want to draw is that the idealised nature of the stipulations that underlie our semantic theory makes many semantic expla-

¹⁶The case of gradable adjectives is special in that context-sensitive object expressions from the object language are used in the metalanguage only in a limited way. (See Glanzberg (2014).) But the point could be made with other sorts of example: for example, on almost any reasonable semantics, all of the claims involved in the explanation of why “Ansel has brown hair” entails “Ansel has hair” would be true on every interpretation of “brown” and “hair”.

nations possible *even if we are willing to grant the radical contextualist almost everything she wants*.

It would be possible to rest content with this conclusion. But in my view the semanticist should want more. The general phenomenon we have seen is that systems of measurement become more useful as the fixed points become more precise. (A platinum-iridium bar is better than the king's foot; the distance traveled by light in a precisely specified interval is better still.) If nothing else, the radical contextualist has observed an interesting linguistic phenomenon that deserves theoretical attention. And while a lexical entry like (10) might be good for many things, it is not very useful for explaining what is going on in the case of the painted leaves.

Why not? To a first approximation, the problem is that the semantic value of "green" stipulated in (10) does not distinguish different uses of "green". Some contexts are such that painted leaves count as "green", and some are not; but an idealised stipulation of (10) provides the same semantic value in either case. So it looks like we just don't have the resources to explain the difference between the two cases.

Let me begin with a concession: this is a genuine explanatory deficit in the theory – something that one would want a theory to explain, but that the theories we are currently considering do not. But this does not in the slightest undermine the explanatory work that our current theories can do successfully. Just as the scale of measuring lengths introduced by (13) is useful for some purposes, and can enter into some explanations, while being inappropriate for other purposes, the system of semantic values introduced in part by (10) can be useful for many things, even though there are things that it cannot explain.

So the objection to semantics cannot be that it is useless – only that it is incomplete (which is something no one serious would doubt). The question, then, is how to improve it. Now in the case of length, our scales were improved by finding more precise fixed points; as we moved from the king's foot, to the platinum-iridium bar, to the speed of light, our system became useful for more and more. I suggest that the same is true for semantic theory. Our strategy for moving forward with respect to the radical contextualist cases should be to move to system introduced by less idealised stipulations.

The problem with idealisation as we are conceiving of it is that an idealised theory does not distinguish cases that may need to be distinguished for some descriptive or explanatory purpose. On our supervaluational model, for example, the claim that "This tree is exactly 3,000 feet tall" and the claim that "This tree is exactly 3,001 feet tall" will both be truth-valueless; so the king's-foot standard will not suit any purpose that requires making such fine

distinctions in length. If we have such purposes, the right strategy is to work to reduce the required idealisation by introducing units fixed by stipulations that are closer to genuine fixed points.

Similarly, given stipulations like (10), utterances of “The leaves are green” will be assigned the same semantic value, regardless of whether the speakers and their audience are interested in photography or in biochemistry. If we have purposes that require distinguishing these cases, the right strategy is to work to reduce the required idealisation. In the case of semantics, the most likely way of realising this strategy is to revisit the stipulations that fix the semantic values of atomic expressions. Consider again Szabo’s proposal: the semantic value of “green” will be something like (16), where P is a contextually supplied specification of the part of a thing that must be green:¹⁷

$$(16) \quad \llbracket \text{green} \rrbracket = [\lambda P. \lambda x. \text{the } P \text{ part of } x \text{ is green}]$$

If our semantic theory uses a stipulation like (16) rather than (10), then we need not assign the same semantic value to all uses of “green”, so that we can distinguish the utterance of “The leaves are green” to the photographer from the utterance of the same sentence to the botanist.

As I have already pointed out, this proposal does not eliminate all relevant context-sensitivity (even with respect to colour adjectives). But that again is a mark of incompleteness rather than failure; our strategy does not rely on the thought that all context sensitivity is eliminable. In any domain, it is an empirical question whether there are genuine fixed points. It could have turned out, for example, that there are no phenomena that always take place at an absolutely fixed temperature. But this would not have made scales fixed in terms of the freezing point and steam point of water useless; nor would it have made futile attempts to isolate factors that affect the freezing and steam points of water, and so to home in on more precise fixed points. (One can eliminate some sources of variation without hoping to eliminate all such sources.) The process is just the kind of bootstrapping described in the previous section: we begin with idealised stipulations, and this makes possible improved theories that enable us in turn to improve our stipulations.

8 Conclusion

We have sought an *account of the nature of semantic theorising* that makes sense of the fact that semantics has made progress despite the radical contextualist

¹⁷I am abstracting away from the fact that Szabo thinks (plausibly) that colour adjectives are gradable.

objections, while meeting several additional desiderata:

Metasemantic Desideratum An account of the nature of semantic theorising must make possible an account of how expressions come to have as their semantic values sets, functions, and similar abstracta.

Proliferation Desideratum An account of the nature of semantic theorising must explain (or explain away) the apparent fact that there are many systems of abstracta that could do the same explanatory work.

Target Ignorance Desideratum An account of the nature of semantic theorising should make it possible to understand how we have made progress in semantic theorising despite radical disagreement and ignorance on the target of that theorising.

On the view we have developed, the Metasemantic Desideratum and the Proliferation Desideratum are met by the claim that semantic values are assigned by theorists setting up a system for representing facts about meaning; they are a matter of stipulation, and different stipulations can be equally good for a given representational purpose. The Target Ignorance Desideratum is met by the observation that a system of measurement can be set up using idealised stipulation even by theorists who are relatively ignorant about the nature of the phenomena being measured. The idealisation that stipulation can involve also gives us the key to understanding how semantic progress is possible in spite of the radical contextualist phenomena. And it shows how further progress in semantics is possible: by further theorising, even while admitting that our theories may never be complete.¹⁸

References

Benacerraf, P. (1983). What numbers could not be. In Benacerraf, P. and Putnam, H., editors, *Philosophy of mathematics: Selected readings, second edition*, pages 272–294. Cambridge University Press, Cambridge.

Borg, E. (2004). *Minimal Semantics*. Oxford University Press, Oxford.

¹⁸I am indebted to audiences in Dubrovnik, at the Australian National University, and at the Universities of Bochum, East Anglia, Edinburgh, and St Andrews, and to comments from and conversations with Emma Borg, Alex Byrne, Yuri Cath, Jonathan Cohen, John Collins, Andy Egan, Miguel Egler, Ephraim Glick, Dirk Kindermann, Kevin Scharp, and Ravi Thakral. I am especially indebted to a penetrating question from Paul Pietroski, an illuminating conversation with Hanti Lin, and a helpful email exchange with Bruno Jacinto.

- Burgess, A. and Sherman, B., editors (2014). *Metasemantics*. Oxford University Press, New York.
- Cappelen, H. and Lepore, E. (2005). *Insensitive Semantics: A Defense of Semantic Minimalism and Speech Act Pluralism*. Blackwell, Oxford.
- Chang, H. (2004). *Inventing Temperature: Measurement and Scientific Progress*. Oxford University Press, Oxford.
- Chomsky, N. (1995). Language and nature. *Mind*, 104:1–61.
- Davidson, D. (2001a). The emergence of thought. In *Subjective, Intersubjective, Objective*, pages 123–134. Oxford University Press, Oxford.
- Davidson, D. (2001b). Truth and meaning. In *Inquiries into Truth and Interpretation, second edition*, pages 17–42. Oxford University Press, Oxford.
- Ellis, B. (1968). *Basic Concepts of Measurement*. Cambridge University Press, Cambridge.
- Glanzberg, M. (2014). Explanation and partiality in semantic theory. In Burgess and Sherman (2014).
- Heim, I. and Kratzer, A. (1998). *Semantics in Generative Grammar*. Blackwell, Oxford.
- Jacobson, P. (2002). The (dis)organization of the grammar: 25 years. *Linguistics and Philosophy*, 25:601–626.
- Kaplan, D. (1977). Demonstratives: An essay on the semantics, logic, metaphysics, and epistemology of demonstratives and other indexicals. In Almog, J., Perry, J., and Wettstein, H., editors, *Themes From Kaplan*, pages 481–564. Oxford University Press, New York.
- Kennedy, C. (1997). *Projecting the adjective: The syntax and semantics of gradability and comparison*. PhD thesis, University of California - Santa Cruz.
- Kennedy, C. (1999). Gradable adjectives denote measure functions, not partial functions. *Studies in the Linguistic Sciences*, 29.
- Kennedy, C. (2001). Polar opposition and the ontology of ‘degrees’. *Linguistics and Philosophy*, 24:22–70.
- Kennedy, C. and McNally, L. (2005). Scale structure, degree modification, and the semantics of gradable predicates. *Language*, 81:345–381.

- King, J. C. (2014). The metasemantics of contextual sensitivity. In Burgess and Sherman (2014), pages 97–118.
- Larson, R. and Segal, G. (1995). *Knowledge of Language*. MIT Press, Cambridge, MA.
- Lepore, E. (1983). What model-theoretic semantics cannot do. *Synthese*, 54:167–187.
- Lepore, E. and Ludwig, K. (2007). *Donald Davidson's Truth-Theoretic Semantics*. Oxford University Press, Oxford.
- Lewis, D. (1970). General semantics. *Synthese*, 22:18–67.
- Lewis, D. (1983). Languages and language. In *Philosophical Papers Volume I*. Oxford University Press, New York.
- Matthews, R. J. (2007). *The Measure of Mind: Propositional Attitudes and Their Attribution*. Oxford University Press, Oxford.
- Pietroski, P. (2005). *Events and Semantic Architecture*. Oxford University Press, Oxford.
- Scharp, K. (2013). *Replacing Truth*. Oxford University Press, Oxford.
- Scharp, K. (ms.). On the indeterminacy of the meter.
- Soames, S. (1989). Semantics and semantic competence. *Philosophical Perspectives*, 3:575–596.
- Stalnaker, R. C. (1984). *Inquiry*. The MIT Press, Cambridge, MA.
- Swoyer, C. (1991). Structural representation and surrogative reasoning. *Synthese*, 87:449–508.
- Szabó, Z. G. (2001). Adjectives in context. In Kenesei, I. and Harnish, R., editors, *Perspectives on Semantics, Pragmatics, and Discourse*. John Benjamins, Amsterdam.
- Tal, E. (2011). How accurate is the standard second? *Philosophy of Science*, 78:1082–1096.
- Travis, C. (2008). Pragmatics. In *Occasion-Sensitivity: Selected Essays*, pages 107–128. Oxford University Press, Oxford.

van Fraassen, B. C. (2008). *Scientific Representation: Paradoxes of Perspective*. Oxford University Press, Oxford.

Yalcin, S. (2014). Semantics and metasemantics in the context of generative grammar. In Burgess and Sherman (2014).