

Semantic Plasticity and Speech Reports

Cian Dorr and John Hawthorne

Near-final version: 31 December 2013

Forthcoming in *Philosophical Review*

1. Introduction

Understand ‘semantic facts’ broadly, so as to include facts about what speech acts people perform as well as facts about the semantic values of linguistic expressions. We will assume that semantic facts supervene, at least nomologically, on facts specifiable in some “lower level” vocabulary, such as that of microphysics.¹ We can imagine two models for how such supervenience might work for a given such fact. On the first model, the fact in question supervenes on the microphysical facts in the same way as the fact that a certain person’s height is between 184 and 185 centimetres. At most worlds where the proposition corresponding to the fact is true, tiny changes in the underlying microphysical facts will not take one to worlds where the proposition is false. On the second model, the proposition is more like the proposition that four particular point-particles are arranged in a perfect square: whenever it is true, there are arbitrarily tiny microphysical differences which would make it false. We will refer to propositions of the first kind as *modally robust*, and those of the second kind as *modally plastic*.

In this paper we will be exploring a certain argument that many semantic propositions are modally plastic. This argument does not apply across the board. In certain special cases, where a content is singled out by “joints in nature”, there is little temptation to deny the modal robustness of semantic propositions that attribute that

¹ While the thesis that semantic facts *metaphysically* supervene on microphysical facts is quite controversial, the thesis of *nomological* supervenience is much less so, although it has recently been disputed by Kearns and Magidor (2012) and Hattiangadi (MS).

content. For example, it is plausible that no small difference in the (current) microphysical facts will take us to a possible world where the proposition that 'positively charged' means *positively charged* is false. But such cases are exceptional. Most of the time, an expression's actual meaning is surrounded by a vast cloud of slight variants, all of which seem equally well qualified to be possible meanings. In these cases, it is hard to resist the conclusion that the truth about the expression's meaning is modally plastic to a high degree: when there is such an abundance of potential meanings available, and no differences between these meanings of the sort which could explain why some of them would be much easier to latch onto than the rest, the selection of one as a certain expression's actual meaning has to be extremely sensitive to the details of microphysics. Similarly, when there are so many potential speech acts a person could have performed, and no differences between those acts of the sort which could plausibly be thought to make some of them much easier to perform than the rest, the fact that a certain act was actually performed has to be extremely sensitive to the details of microphysics. Let us call arguments of this general kind *arguments from abundance*.²

² Our argument here echoes an argument of Williamson's for the claim that in many cases, 'small shifts in use induce correspondingly small shifts in reference': 'And what is the alternative to the postulated correspondence between shifts in use and shifts in reference? The reference of a vague term would be insensitive to most small shifts in use, but would change in large jerks as the pattern of use shifted across critical boundaries. The privileged extensions would mark out something like hidden natural kinds. Situations like the following would obtain. "Bald" actually refers to the property of having fewer than 3,832 hairs on one's scalp. If we had been as likely to apply "bald" to someone with n hairs as we actually are to apply it to someone with $n+100$ hairs (for all n), then "bald" would still have referred to the property of having fewer than 3,832 hairs on one's scalp. But if we had been as likely to apply "bald" to someone with n hairs as we actually are to apply it to someone with $n+1,000$ hairs (for all n), then "bald" would have referred to the property of having fewer than 2,832 hairs on one's scalp. 3,832 and 2,832 are natural boundaries of a sort; 3,732 is not. That is implausible, although not incoherent' (Williamson 1997: 948–9).

It is important to note that this mode of argument in no way depends on Williamson's distinctive view of vagueness. *Whatever* sort of semantic profile one assigns to the word 'bald', it is hard to see how it could fail to belong to a huge set of possible semantic profiles, within which the postulation of 'natural boundaries' would

In thinking about whether we could embrace the conclusion that the truth about a certain semantic subject-matter—say, the meaning of the word ‘salad’—is modally plastic, it can be helpful to consider analogous claims about the way in which the truth about that subject matter varies over time. As regards this subject matter, proponents of *temporal plasticity* will hold that each meaning ‘salad’ has ever had has been had for a very short time—perhaps only for a single instant, while proponents of *temporal robustness* will hold that each meaning ‘salad’ has is had for quite a while.³ While temporal plasticity can provide a helpful analogy with modal plasticity, it is important to note that it is perfectly consistent to posit widespread modal plasticity in the semantic facts while rejecting temporal plasticity. Even if the semantic history of ‘salad’ involves only a few meanings, each of which prevails for many decades, the identity of the sequence of meanings could be exquisitely sensitive to certain details of microphysical history. There is no obvious way in which such a position would fail to do justice to the motivating thought behind the argument from abundance.⁴

Although the view that temporal plasticity is widespread in the semantic domain is radical, and not clearly motivated by the arguments from abundance which motivate

be no more plausible than the postulation of such boundaries within the set of properties of the form *having fewer than n hairs on one's head*. Williamson's positive theory of vagueness does suggest a very different argument for widespread modal plasticity, as follows: most words are vague; vagueness is due to modal plasticity; so most facts about word-meaning are modally plastic. We will not be discussing this argument (but see Hawthorne 2006, ch. 9).

³ It is less obvious what the temporal analogue would be of the claim that certain propositions about speech acts are modally plastic. The least distracting option is to focus on the number of instances of a given speech-act type: for example, the proponent of temporal plasticity will hold that in the entire history of the world there are only a few acts that are assertions that salad is delicious (perhaps only one, or none at all).

⁴ Is there also a coherent position that endorses temporal plasticity while denying modal plasticity? It is certainly possible in principle that although an expression has any given meaning only for an instant, each meaning it ever has is had *at some time or other* throughout a wide range of nearby worlds. But, while there is something to be said for this view in certain special cases (where, intuitively, there is only one dimension along which meanings can vary), we find it hard to motivate this as a general model.

modal plasticity, there is a different kind of argument for widespread temporal plasticity which is not without its attractions. Suppose—not implausibly—that the use of ‘salad’ has evolved gradually over the last few centuries, from a situation in which people only applied the word to cold dishes with a high preponderance of green leaves of some sort, to the current state where people comfortably apply the word to various warm, leaf-free concoctions. The meaning of ‘salad’ is evidently not the same at the beginning and at the end of this history. It follows that there is at least one second-long interval such that the meaning of ‘salad’ is not constant throughout that interval. But given the gradual character of the underlying changes, all the second-long intervals are extremely similar in relevant respects. This makes it tempting to conclude that semantic changes occur during *all* of them, or at least during some non-trivial proportion of them.

Call arguments like this one *arbitrary-boundaries arguments*. Similar arguments can be given on behalf of modal plasticity, in which we consider a long series of possible worlds, each only slightly different from its neighbours in microphysical respects, stretching from the actual world to one in which the use of ‘salad’ is very different, and conclude from the fact that there are no non-arbitrary grounds for positing semantic discontinuities at just a few points in the series, all or most of the worlds in the series differ from their neighbours as regards the meaning of ‘salad’. Whatever one might make of these arbitrary-boundaries arguments, it is important to distinguish them from arguments from abundance, which are our primary focus. Arguments from abundance turn on a claim of parity as regards candidate *meanings*, whereas arbitrary-boundary arguments turn on a claim of parity as regards candidate *locations for semantic boundaries*. It would not be inconsistent to accept the former but not the latter.

There is much that cries out for clarification and justification in the arguments for modal and temporal semantic plasticity we have just been considering; in §2, we will

make the argument from abundance a lot more rigorous. But before we turn to this task, let us consider a *prima facie* compelling argument against both temporal and modal plasticity, based on our practices of reporting speeches made at other times and other possible worlds. We will begin with the temporal case, where this case is most clear-cut. In a vast array of circumstances, we confidently use homophonic methods in reporting speeches made in the not-too-distant past. For example, we hear Sally saying 'Salad is delicious', and five minutes later we utter

- (1) Sally asserted that salad is delicious.⁵

This ordinary practice is hard to square with the idea that the meanings of words like 'salad' are in constant flux. If 'Salad is delicious' meant something different five minutes ago from what it means now, it seems that we will be saying something false in uttering (1). For presumably the proposition Sally asserted was the one semantically expressed by 'Salad is delicious' at the time of her speech, while the proposition we characterised her as having asserted in uttering (1) was the distinct proposition semantically expressed by 'Salad is delicious' at the time of *our* speech.⁶

Modal plasticity threatens to generate analogous problems with counterfactuals whose consequents involve speech-act reports, such as (2)–(4):

- (2) If we had taken Sally to Giorgio's, she would have asserted that salad is delicious.
- (3) If I had the courage, I would tell them to stop making fun of my baldness.

⁵ We will often use *assert* in our examples rather than more colloquial verbs like *say* and *tell*: this is because of *assert*'s convenient syntactic flexibility ('They asserted the same proposition' is grammatical while 'They said the same proposition' is not). We stress that we are not using 'assert' in any special technical sense.

⁶ Kearns and Magidor (2008) make essentially this argument against temporal plasticity.

- (4) No matter which graduate school I had gone to, I would probably still have been insisting that monetary policy can prevent recessions.

There seem to be many settings where we could reasonably make such claims with confidence. But this is hard to square with the claim that we should regard facts about the meanings of words like 'salad', 'baldness', and 'recession' as modally plastic. After all, many microphysical facts would have to be different for the antecedents of (2)–(4) to be true. If the truth about the meanings of these words is exquisitely sensitive to certain microphysical parameters, on what grounds could we be confident that these parameters would still have had their actual values in the relevant counterfactual worlds? But if we thought it unlikely that 'Salad is delicious' would have *meant* that salad is delicious if we had taken Sally to Giorgio's, then it is hard to see how we could be confident that Sally would in those circumstances have *asserted* that salad is delicious, rather than asserting the slightly different proposition that 'Salad is delicious' would have semantically expressed.⁷

Thus we are faced with a paradox, or at least a choice between uncomfortable options. On the one hand, the view that typical semantic facts are modally plastic seems to be required by the plethora of possible contents and the lack of any basis for regarding some of them as much easier to latch onto than others. On the other hand,

⁷ It is arguable that analogous issues arise for cross-temporal and counterfactual attitude reports, like 'Five minutes ago, Sally wanted to eat a salad' and 'If I hadn't gone to Harvard, I would have believed that some truths are analytic'. One way of developing these problems appeals to the not uncommon view that *believing that P* and *wanting that P* require standing in certain special relations to a sentence in one's language of thought that semantically expresses the proposition that P. This allows one to argue from temporal or modal plasticity in facts about language-of-thought semantics to the falsity of typical cross-temporal and counterfactual attitude reports. We can also construct an argument from abundance which applies directly to attitude reports, without needing to assume the language-of-thought model. But let us postpone thinking about this until we have a more rigorously formulated argument from abundance on the table (in §2).

such widespread plasticity makes it mysterious how it could ever be reasonable to accept counterfactual speech reports like (2)–(4).

This puzzle is reminiscent of other, more familiar arguments about speech reports. Semantic externalists often complain that semantic internalism predicts that we commonly mean different things by our words, and hence that our reports of one another's speeches are systematically false (Burge 1979). Similarly, anti-contextualists about a particular word often complain that contextualism about that word predicts that it rarely if ever happens that two utterances of the word contribute exactly the same content, so that speech reports involving that word are far more precarious than we normally suppose (Cappelen and Lepore 2005). What has not been fully appreciated is that, once we turn our attention to cross-temporal and counterfactual reporting, analogous problems arise for nearly everyone.⁸

Our goal in the rest of the paper is to give a more rigorous presentation of the puzzle and to explore several strategies for resolving it. §2, as already advertised, will

⁸ One especially close set of problems involves cross-linguistic speech reports, even holding time and world fixed. Might the considerations that make trouble for cross-time and counterfactual speech reports also carry over to our reports of assertions by French speakers? There are indeed highly stylised cases where an analogue of our arbitrary-boundaries argument for temporal plasticity in the meaning of 'salad' (in a world of gradual change from one usage to a very different usage) can be made for synchronic reports. Imagine a series of villages strewn along the bank of a long river, where each village is in communication with its immediate neighbours. At one end, people use 'salad' like us; at the other end, they apply it only to leafy dishes, but the differences between neighbouring villages are small. There is some temptation to say that 'salad' means something different in each village, on the grounds that no pair of adjacent villages has a better claim to differ semantically than any other. However, even if we do say this, we shouldn't be too quick to think that it raises any problems for our actual reports of speeches currently being made in, e.g., French. Those who want to save cross-linguistic reports in the actual world have many resources: for example, they could claim that the existence of established translational practices explains the semantic uniformity of 'salad' and 'salade' (given the absence of the countervailing pressures that force non-uniformity in the river world). Note that for similar reasons, even if we granted that the meaning of 'salad' is subject to temporal plasticity given our toy assumptions about its history, we would have no straightforward grounds for positing the same kind of temporal plasticity in worlds where the gradual changes in the use of 'salad' are confined within modest boundaries.

sharpen up the intuitive notion of modal plasticity we have been working with so far and give a rigorous reconstruction of the argument from abundance, in the form of a classically valid argument-schema with two premises. §3 considers whether the conclusion of this argument can be reconciled with ordinary counterfactual speech reports like (2)–(4). §4 and §5 discuss views that deny the argument's first and second premise, respectively. (The views discussed in §4 involve saying that, among the abundance of seemingly similar candidate meanings, a few are in fact much easier to latch onto than the rest, while those discussed in §5 involve the claim that we often say infinitely many things by uttering a single sentence). §6 concludes. Our survey will not include views that escape the puzzle by rejecting theorems of classical logic, such as instances of the law of the excluded middle: we leave it to the advocates of such revisions to investigate whether they can provide an attractive alternative to the options we will consider.

2. The case for plasticity

For some words—perhaps including 'electron', 'spacetime point', 'plus', 'identical', 'something'...—there is no compelling reason to posit semantic variation at close worlds. What is the relevant difference between these words and words like 'salad', which makes it plausible to think that the latter have a different sort of trans-world semantic profile? The basic thought that drives the argument from abundance is that in the case of 'salad', there is a large family of properties that are *candidates* to be expressed, all of which are roughly alike in the respects that seem to matter for the expression relation. Because of these similarities, it is hard to believe that there are big differences among the modal profiles of the propositions which ascribe to these candidate meanings the property of being semantically expressed by 'salad'. In

particular, it is hard to believe that some of the candidates get expressed at a big chunk of modal space, while others only get expressed in a miniscule sliver, or nowhere at all. But then, since there are so many candidates, and since only a few of them get expressed at any given world, the set of worlds where any given candidate is expressed must be tiny.⁹

To make further progress with this somewhat more careful version of the argument from abundance, we will need to make sense of the rather mysterious notion of size for sets of worlds. Contrary to what the metaphor of 'modal space' might suggest, it is unclear whether there is any pre-theoretically intelligible measure on the set of all metaphysically possible worlds. Fortunately, the argument requires nothing so general: to generate a puzzle about ordinary counterfactual speech reports, we may not need to look beyond the sphere of nomologically possible worlds, and in that case we can achieve some more discipline by borrowing some tools from physics. Classical statistical mechanics makes extensive use of a measure (the Liouville measure) over *phase space*—the set of possible instantaneous physical states of the world in which certain basic unchangeable parameters (for example, the number of particles and their masses) are held fixed. Assuming that the laws are classical and fully deterministic, and that all truths supervene nomically on the truth about the history of instantaneous

⁹ It is tempting to invoke facts about the naturalness of properties (in the sense of Lewis 1983) to explain the difference between cases where semantic plasticity is *prima facie* plausible and cases where it is not. But one must be careful in spelling out this connection: the argument from abundance works even for words that express perfectly natural properties, so long as those properties belong to large families of similar perfectly natural properties. (For example, Lewis seems to think that all of the continuum many determinate distance relations are perfectly natural.) More generally, there are many reasons to doubt that any ranking of properties plays all the roles that Lewis wanted the naturalness-ranking to play, including connections to supervenience, duplication, lawhood, resemblance, and induction as well as psychology and semantics (see Dorr and Hawthorne forthcoming). This generates the potential for there to be multiple useful notions of naturalness, satisfying different consistent fragments of the full Lewisian role; these notions might be related to semantic plasticity in quite different ways.

physical states, there is a straightforward way of transferring the Liouville measure from the set of phase-space points to the set of “physically accessible” worlds—nomically possible worlds which match actuality as regards the basic unchangeable parameters such as the number of particles.¹⁰

In addition to the notion of size as absolute phase-space measure, the classical framework also allows for further fine-grained distinctions. For a wide range of ‘well-behaved’ sets S in phase-space, there is a mathematically natural notion of the measure *relative to S* of a subset R .¹¹ And as before, this notion of relative measure extends in a natural way from sets of phase-space points to sets of possible worlds. These relative measures are especially interesting when S itself has phase-space measure zero. For example, on standard physical assumptions any *energy hypersurface*—any maximal set

¹⁰ This fact flows from *Liouville’s theorem* which says that the Liouville measure is invariant under temporal evolution. In the simplest case, each phase space point p and duration Δt will determine points $p_{+\Delta t}$ and $p_{-\Delta t}$, such that it is nomically necessary that if the universe is ever in state p , it will be in state $p_{+\Delta t}$ Δt later and was in state $p_{-\Delta t}$ Δt earlier. In this setting, Liouville’s theorem entails that for any Δt and measurable set S , $\{p_{+\Delta t}: p \in S\}$ and $\{p_{-\Delta t}: p \in S\}$ have the same measure as S itself. Let us count facts about which instants of time there are and about the temporal distances between instants among the basic parameters upon which all physically accessible worlds agree. Given this, it follows from the previous result that for any set of physically accessible worlds W and instants t and t' , the set of phase space points instantiated at t at worlds in W and the set of phase space points instantiated at t' at worlds in W have the same measure (if either is measurable). We can thus define our measure on the physically accessible worlds by identifying the measure of W with the measure of the set of phase-space points instantiated at worlds in W at any arbitrarily chosen instant. (A more complicated version of this method can be made to work even if the phase space contains points for which the time-evolution operation is ill-defined for certain values of Δt , so long as we stipulate that all particles permanently cease to exist or come into existence for the first time whenever the time-evolution is ill-defined.)

This definition requires making sense of instants of time as entities existing in multiple possible worlds. Those who are suspicious of this reification of times will presumably think that nomologically possible worlds never differ merely by a temporal shift. Since on such a view, no phase-space point is instantiated at more than one nomologically possible world, we can simply identify the measure of a set of worlds with the measure of the set of phase-space points instantiated at any time at any world in the set.

¹¹ The definition of this relative measure is straightforward whenever S is an *embedded symplectic submanifold* of phase space; more complex techniques can be used to define natural relative measures in certain other cases.

of phase-space points which agree as regards the total energy of the system—has phase-space measure zero, but will generate well-behaved relative measures.

Of course, the physics of the real world is not classical. But measures on sets of worlds that would suit our purposes also arise naturally on all known interpretations of quantum physics. For the sake of simplicity and concreteness, let us pretend from now on that the actual world is a classical world with finitely many particles (so that we do not have to deal with the mathematical complexities introduced by infinite-dimensional phase-spaces), moving in a fixed background spacetime in accordance with fully deterministic force laws.¹²

Given this framework, we can use the phase-space measure to articulate various more precise versions of the claim that certain propositions hold at ‘tiny’ sets of worlds. For the sake of keeping the argumentation as clean as possible, we will focus in particular on interpretations where ‘tiny’ sets are those with *zero* phase space measure, either absolutely or relative to some physically natural set of worlds, such as the set of worlds whose total energy is the same as that of the actual world. As we will see, a strong argument can be made that the sets of worlds corresponding to various semantic propositions are ‘tiny’ even in this very demanding sense.

The notions of size provided by physics suggest one avenue for quantifying the informal notion of modal plasticity, or “degree of dependence on the exact details of microphysics”, which we introduced in §1. But there are also other reasonable avenues for explicating that notion. Notably, there are possible interpretations that appeal to some notion of *distance* between possible worlds, guided by questions like ‘How far, on

¹² You might worry that such a world could not possibly contain enough stable complexity to undergird the existence of people who perform any speech acts at all. This is surely right as regards worlds where the force laws take a simple form, like those of Newtonian gravitation. But it is plausible we can get there to be people by having recourse to complex, gerrymandered force laws.

average, does one have to go to get from a world where the proposition in question is true to one where it is false?'.¹³ (In the setting of classical physics, there are various reasonably natural ways to define metrics over certain sets of physically possible worlds.) These metric-based ways of quantifying plasticity will come apart from the measure-theoretic approach in many cases. However, sets of worlds with *zero* measure will plausibly also count as maximally plastic according to a metric approach. For on any physically reasonable notion of distance between worlds, each 'sphere' comprising all worlds within some positive distance ε of a given world will have positive phase-space measure. Thus no zero-measure set can contain any such sphere—in other words, every zero-measure set must be *nowhere dense*. And being nowhere-dense should suffice for being maximally plastic on a metric approach, since each world in a nowhere-dense set is arbitrarily close to worlds not in that set.¹⁴ Thus we need not

¹³ 'On average' here could perhaps be cashed out using the measure-theoretic notion of expected value.

¹⁴ The converse inference fails, as we can see by considering, for example, the set of worlds where the distance between particles a and b at t is an irrational multiple of the distance between particles c and d at t . Given minimal assumptions about the relevant notion of distance, this set will be nowhere dense, and hence presumably maximally plastic on a metric interpretation of plasticity. However, being the complement of a zero-measure set, it is very far from being maximally plastic on a measure-theoretic interpretation. Examples like this also suggest that the metric notion of plasticity is less useful than the measure-theoretic one for the purpose of generating a puzzle about our right to be confident in counterfactuals whose conclusions are modally plastic. For we would need to rely on some quite controversial claims about counterfactuals to justify the claim that we cannot reasonably be confident that (for example) if Romney had been elected in 2012, the ratio of the distance between a and b to the distance between c and d would have been irrational. On Lewis's analysis (Lewis 1973), according to which the truth of the counterfactual requires the distance-ratio to be irrational throughout a non-empty set of world where Romney wins all of which are more similar to actuality than any worlds where Romney wins and the ratio is rational, there is strong reason to regard the counterfactual as false. But on other accounts (e.g. Moss 2013), the counterfactual will deserve the same high level of credence as the proposition that the relevant distance-ratio is in fact irrational, for the same reason that we should assign credence $\frac{1}{2}$ to the proposition that a certain coin (which we are sure is fair) would have landed heads had it been tossed.

assume a measure-theoretic analysis of modal plasticity to justify taking the claim that certain propositions have zero measure to entail that they are extremely modally plastic.

Having settled on zero measure as our working interpretation of 'tiny', we are in a position to provide a more rigorous and general version of the argument from abundance. What we really have is an abstract argument-schema, for showing that each member of a certain uncountable set of propositions Γ has measure zero relative to a certain set W of physically accessible worlds (assuming it has a well-defined measure relative to W at all).

The continuum argument

Premise 1 (*Parity*): If any member of Γ has positive measure relative to W , uncountably many do.

Premise 2 (*Finitude*): Only finitely many members of Γ are true at any given world in W .

Conclusion (*Plasticity*): No propositions in Γ have positive measure relative to W .

This argument-schema is valid, given that the phase-space measure is ' σ -finite' (that is, every measurable set is the union of countably many finite-measure subsets). Its validity follows from the following mathematical fact: if each point in a σ -finite measure space belongs to only finitely many members of a certain family of sets, then at most countably many members of that family can have positive measure.¹⁵ Pictorially: there

¹⁵ *Proof*: Let the measure be μ and let Γ be the family. Suppose that for some finite-measure set S and $\varepsilon > 0$, there were infinitely many sets $P \in \Gamma$ such that $\mu(P \cap S) \geq \varepsilon$. Then there would be a countable sequence P_0, P_1, P_2, \dots such that $\mu(P_i \cap S) \geq \varepsilon$ for all i . Let $F_i = S \cap (P_i \cup P_{i+1} \cup P_{i+2} \cup \dots)$, and let F be the intersection of all the F_i . Clearly $\mu(S) \geq \mu(F_i) \geq \mu(S \cap P_i) \geq \varepsilon$ for each i , so (by the fact that measures are "continuous from above") $\mu(F) = \lim_{i \rightarrow \infty} \mu(F_i) \geq \varepsilon$. But a point belongs to F if and only if it belongs to infinitely many of the P_i , so by the stipulation that no point belongs to infinitely many

is no way to cover the plane with an uncountable collection of tiles, each with positive area, without some point ending up under infinitely many tiles.¹⁶

Let's look at a *prima facie* compelling instance of the continuum argument. On the table in front of us are two cups, Cup A and Cup B. They are pretty far apart right now. Call the proposition that Cup A and Cup B are pretty far apart at t , which we just asserted, PrettyFar. PrettyFar belongs to a large class of similar propositions about Cups A and B, which differ as regards how far apart they have to be at t for them to be true. Since there are continuum many such distance thresholds, there will be a continuous set Θ of variants of PrettyFar, even if we restrict ourselves to variants that differ only slightly—say, by at most a centimetre. Let Γ contain, for each p in Θ , the proposition

members of Γ , F is empty and thus has measure 0. . So by *reductio*, we can conclude that for every finite-measure S and every $\varepsilon > 0$, there are only finitely many $P \in \Gamma$ for which $\mu(S \cap P) \geq \varepsilon$. This entails for each finite-measure S , the set of $P \in \Gamma$ for which $\mu(S \cap P) > 0$ is countable, since we can enumerate it by first listing those P for which $\mu(S \cap P) \geq 1$, then those for which $\mu(S \cap P) \geq 1/2$, then those for which $\mu(S \cap P) \geq 1/3$, etc. But since the space is σ -finite, there is a countable collection of finite-measure sets S_0, S_1, S_2 that cover the whole space, in which case every positive-measure member of Γ must have a positive-measure overlap with at least one S_i . So the set of positive-measure members of Γ is the union of countably many countable sets, and hence countable.

Note that for the proof to work it is enough to assume that the set of points belonging to infinitely many members of Γ has zero-measure: it doesn't actually need to be empty. Thus the continuum argument would remain valid if we replaced *Finitude* with the following weaker premise: that the set of worlds where infinitely many members of Γ are true has measure zero relative to W .

¹⁶ *Parity* as we have stated it is weaker than the claim that the propositions in Γ have zero measure relative to W , since it allows that some or all members of Γ might be unmeasurable relative to W . We find it implausible that any of the propositions we are interested in are unmeasurable, and will not have much to say about approaches to the puzzle that turn on denials of measurability. However, for those who are more open to this idea than we are, it is worth noting a corollary of the result proved in the previous footnote: if *Finitude* is true, then at most countably many members of Γ can have positive *inner measure*, where a proposition's inner measure is defined as the supremum of the measures of the measurable sets of worlds throughout which it is true. Since the considerations that motivate *Parity* also motivate the claim that if any members of Γ have positive inner measure relative to W , uncountably many do, those who accept the continuum argument for Γ and W are under strong pressure to claim that every member of Γ has zero inner measure relative to W .

that it is asserted at least once within a billion years of t .¹⁷ And let W be the set of physically accessible worlds with the same total energy as the actual world.

The plausibility of *Parity* for this choice of Γ and W arises from the thought that the different distance relations ascribed by members of Θ are equally natural—there are no ‘joints in nature’ between them. Since there are no significant differences in metaphysical status among the members of Θ , it would be strange if a few of the members of Θ differ from the rest in being asserted throughout positive-measure sets of worlds (within a billion years of t). For example, it would be weird if the member of Θ with the truth-condition that Cup A is at least 78.4 centimetres from Cup B was quite a popular target of assertion across the space of physically accessible worlds, while the member with the truth-condition that Cup A is at least 79.3 centimetres from Cup B was asserted almost nowhere. Of course, in principle, one could think that this is true as regards the space of physically accessible worlds as a whole but false as regards worlds with the same total energy as the actual world: assertions of *PrettyFar* might be widespread on this energy hypersurface, while assertions of other members of Θ were similarly widespread on others. But it is tremendously implausible that this particular hypersurface induces such a semantic boost—the supervenience of the semantic on the microphysical is far more intricate and delicate than that.

To deny *Finitude* in this case would be to claim that there are some worlds on our energy hypersurface at which infinitely many propositions about the distance between Cups A and B are asserted within a billion years of t . On standard pictures of what it

¹⁷ This instantiation of the argument requires assuming that instants can be identified across possible worlds. Those who have worries about such identifications should substitute some qualitative description of the relevant time t . Alternatively, in an A-theoretic perspective, one could consider temporal proposition like *someone will assert such-and-such within a billion years*—since a point in phase space corresponds to an instantaneous world-state, a measure applicable to temporal propositions can be read directly off the phase space measure. Even B-theorists might find it more convenient to work with measures applying to properties of instants rather than complete propositions.

takes to assert something, such worlds would have to be very strange indeed. For example, one can imagine that at some point people start talking faster and faster, so that they end up making infinitely many speeches about Cups A and B. But such goings-on are plausibly physically impossible—even setting aside relativistic considerations (see Arntzenius 2006), our vocal cords would fly apart long before the speech-making was finished. (In §5 we will discuss alternative views of assertion that make failures of *Finitude* physically realistic.)

If the continuum argument is sound in this case, it is sound for a vast array of alternative choices of Γ and W . To choose another example more or less at random, let Γ contain, for every distance relation between 1cm and 10m, the proposition that some community has, at a certain instant t , a syntactically simple predicate that semantically expresses that distance relation; and let W be the set of physically possible worlds which exactly match the actual world as regards the positions and velocities at t of all particles more than one light year from the Sun.

For this instantiation of the continuum argument, a case for *Finitude* can be made based on the safe assumption that at each world in W , there are only finitely many linguistic communities at t each with only finitely many syntactically simple predicates, and the orthodox semantic assumption (which §5 will scrutinize) that a simple predicate expresses at most one relation in the language of a given community at a given time. (Similar considerations provide *prima facie* support to *Finitude* for a very wide range of choices of Γ and W .¹⁸) The case for *Parity* is as before: there aren't any metaphysical contrasts within the set of distance relations that could plausibly be thought to render

¹⁸ Many philosophers will also be attracted to *Finitude* not just for propositions about meanings and speech acts, but also for propositions about belief, desire, confidence, and other psychological states. While the idea that creatures like us do not have the capacity to believe or entertain infinitely many propositions at once could be argued for on the basis of a “language of thought” account (see note 7), the appeal of this idea is by no means restricted to those who accept such an account.

some of them much more popular than others as semantic values for simple predicates.¹⁹ You might be tempted to think that it is much more common (within W) for predicates to express *rational* multiples of relations like *one metre apart* or *one foot apart* than to express distance relations not simply definable in familiar units—after all, we in the actual world seem have a much easier time latching on to the former than the latter. But this ignores the extreme sensitivity to microphysical detail in the process that led us to choose metres and feet as units of length. Suppose that ‘metre’ was introduced by the reference-fixing description ‘the length of that Parisian bar at t^* ’. The set of worlds where the bar has *exactly* its actual length at t^* has measure zero relative to W . (Remember that the bar’s constituent particles are constantly vibrating.) So even if the practice of using that reference fixing description extended over a finite-measure subset of W , the result is almost always for ‘metre’ to end up expressing some other length. Indeed, it is only at a measure-zero region that the distance *one metre* will even be finitely expressible in terms of the resulting system of units.²⁰

¹⁹ Indeed, if you were looking for a physically tractable measure of ‘eligibility’ in something like Lewis’s sense, you might consider identifying the eligibility of a property with (some function of) the phase-space measure of the proposition that some community expresses it with a syntactically simple predicate.

²⁰ In real-world physics there turns out to be an especially natural length-property, the Planck length. This introduces asymmetries among length-properties that might reasonably be thought to make for differences in the ease with which they can be expressed. For example, one might suggest that rational multiples of the Planck length are expressed by simple predicates at nonzero-measure sets of worlds, whereas almost all irrational multiples of the Planck length are not. Since there is a Planck unit for each physical dimension, real-world physics may not provide any examples of the homogenous, continuous families of contents which make *Parity* look compelling.

Certain views in metaphysics also suggest reasons to doubt the existence of such families. Some philosophers hold that the totality of length-facts supervene on facts about some finite collection of relations, such as the betweenness and congruence relations of Tarskian geometry (see Field 1983, Arntzenius and Dorr 2012, Dasgupta forthcoming). The question how long something is in metres is settled, doubtless in a messy and holistic way, by its role in the pattern of relational facts. While this view does not threaten the thought that the property of being one metre long belongs to an enormous set of properties that are semantically on a par, it does not clearly support the claim that any such set is *uncountable*.

As a final example, let us turn to 'salad'. In this case there are several possible strategies for generating an uncountable set of propositions that includes the proposition that 'salad' means *salad*, and for which both premises of the continuum argument are plausible. We could consider a set of variant meanings that differ along some relevant continuous parameter such as minimal required proportion of plant matter, maximum temperature, or minimum volume. Or, more generally, we could consider any continuous series of cases ranging from some paradigm salad to a paradigm non-salad, and consider some set of properties each of which generates a different cutoff point in this series and is otherwise as similar as possible to the property of being a salad.²¹ If the physical world is continuous, then similar strategies will work for most ordinary predicates.

However, even if one rejects *Parity* on these physical or metaphysical grounds, one will still have to contend with arguments from abundance in which *Parity* is replaced by the claim that if any member of Γ has a *large* measure relative to W , *many* do; *Finitude* is replaced by the claim that *few* members of Γ are true at any world in W ; and *Plasticity* by the claim that no member of Γ has a *large* measure relative to W . Given appropriate precisifications of 'large', 'many' and 'few', we can get this argument-form to be valid. When Γ is enormous (even if finite), the relevant sense of 'large' will be very weak, though not as weak as 'positive measure'. Thus, since our actual units are not close to any simple multiples of the Planck units and not simply definable in terms of relations like betweenness and congruence, arguments of this form can be used to support the claim that typical speech reports and meaning attributions express propositions whose measure is *very small* (especially when the relevant contents can be varied along many different dimensions). This weakening of *Plasticity* also poses a *prima facie* threat to our ordinary practice of counterfactual speech reporting, though there are some strategies for reconciliation that have some prospect of success for very small positive measure but are non-starters for zero measure: see footnote 29 below.

²¹ We are thinking here of meanings as determining extensions, whereby any given case is either in or not in the extension of a given candidate meaning. Some philosophers will prefer a different ideology, for example one in which predicate-meanings impose tripartite divisions in to an extension, an anti-extension and a 'zone of undefinedness', or one on which predicate-meanings correspond to functions from cases to 'degrees of truth'. It is equally easy to see how continuous families of variant meanings can be constructed within these alternative semantic frameworks. Of course, one might think that considerations related to vagueness somehow reduce the plausibility of the premises of the instances of the continuum argument generated in this way. We will consider the bearing of vagueness on the case for *Parity* in section 4.3.

We won't try to state a precise condition under which considerations about arbitrariness provide *prima facie* support to *Parity*, for a given Γ and W . But our general thought is that *Parity* will be hard to resist whenever (i) Γ is an uncountable set of propositions each of which characterises an expression as having a particular meaning, or a person as standing in some speech-act relation to a particular proposition; (ii) the set of contents ascribed by members of Γ consists of close variants that are roughly alike in respects that matter to physicists and metaphysicians, and (iii) W is topologically connected, delimited in a way that is reasonably natural from a physical point of view, and big enough to make it plausible that more than one member of Γ is true somewhere within W .

Given the validity of the continuum argument, we have three choices, when faced with some uncountable family of propositions about some semantic subject matter. We could embrace *Plasticity*, and somehow defuse the worry that this leads to a disastrously pervasive rejection of counterfactual speech reports. We could deny *Parity*, and somehow find a way to make our peace with concerns about metaphysical arbitrariness. Or we could deny *Finitude*, and somehow accustom ourselves to the idea that one can perform infinitely many speech acts in uttering a single sentence. §3, §4, and §5 will consider these options in turn.²²

²² In §3 we will focus on views according to which the propositions expressed by typical meaning attributions and speech reports have measure zero. One could also endorse *Plasticity* while claiming that these propositions are typically *unmeasurable*. But we find this view hard to believe. The sets of propositions Γ that feature in the interesting instances of the continuum argument carry a natural topology, inherited from the similarity structure of the set of contents they attribute (in the first two examples discussed above, this is based on the topology of the set of length relations). This topological structure on Γ determines a canonical topology on the set of n -membered subsets of Γ , for each n (viz. the quotient of the product topology by the group of permutations). The natural vision on which the continuum argument is sound is one where small differences in the microphysical facts generate correspondingly small differences in the semantic facts, so that the function from physically accessible worlds where exactly n members of Γ are true to the n -membered subsets of Γ that are true at

First, however, it is worth pausing to consider to what extent an analogue of the continuum argument could be used to support claims of *temporal* plasticity. The obvious way to construct such an argument is to let Γ be a set of *properties of instants* (or temporal propositions) rather than a set of eternal propositions. The analogue of *Parity* will then say that if any member of Γ has positive duration (i.e. has a set of instances with positive temporal measure), uncountably many do; the analogue of *Finitude* will say that only finitely many members of Γ are true at any instant; and the analogue of *Plasticity* will say that no member of Γ has positive duration. However, the case for the temporal version of *Parity* will, typically, be much weaker than the case for the modal version. Consider a view that endorses modal but not temporal plasticity for propositions about which distance-relations are semantically expressed by simple predicates. In the history of each possible world, only a few distance-relations are ever expressed by simple predicates, although for each distance-relation, the set of physically accessible worlds where it is *ever* expressed by a simple predicate has measure zero.

them is *continuous* (using the phase-space topology on the worlds and the similarity-based topology on the subsets). And this continuity claim (together with *Finitude*) entails that for each $p \in \Gamma$, the set of worlds where p is true is measurable. For each n the set of n -membered subsets of Γ that contain p is *closed* according to the similarity-based topology on the set of n -membered subsets of Γ . So its inverse image—namely, the set of worlds where p and exactly $n-1$ other members of Γ are true—is closed in the phase-space topology, and hence measurable. Moreover, by *Finitude*, the set of worlds where p is true along with infinitely many other members of Γ is also measurable, since it has measure zero. Hence the set of worlds where p is true is measurable, being the union of countably many measurable sets.

Moreover, even if one held that ordinary speech reports were unmeasurable rather than zero-measure, one would still face worries about counterfactual speech reports akin to those to be discussed in §3. For while the epistemology of counterfactuals with unmeasurable consequents is admittedly obscure, many reasons for worrying about counterfactuals with zero-measure consequents carry over to counterfactuals whose consequents are unmeasurable but have zero *inner measure*. And as we noted in footnote 16, those who endorse the continuum argument are under strong pressure to say that the relevant propositions have zero inner measure. For these reasons, we think it is implausible that failures of measurability provide the way out of our puzzle, so from now on we will take it for granted that all the propositions of interest have well-defined measures.

There is no obvious way in which this picture fails to respect the idea that the distance relations are 'metaphysically on a par'. The necessary truths about the way semantics supervenes on microphysics treat the distance relations democratically—each gets to be expressed at a measure-zero set of worlds. The only “singling out” is done by nomologically contingent truths—a body of truths which, given determinism, is nomically equivalent to an exact specification of the microphysical state at some time early in the history of the universe. It is much harder to motivate the idea that *this* singling out involves a metaphysically objectionable kind of arbitrariness. We are used to the idea that the contingent truth about actuality might privilege some members of a family of items that are otherwise metaphysically on a par—for example, some mathematically uninteresting number might enjoy the contingent distinction of being the total number of particles.

There is, however, a quite different kind of concern about arbitrariness that can be used to bridge the gap from *Finitude* to *Plasticity* in both the temporal and modal cases, by means of what §1 called an *arbitrary-boundaries argument*. Suppose that a certain set Γ of semantic properties of instants satisfies *Finitude*, so that (by the mathematical fact that underlies the validity of the continuum argument) at most countably many members of Γ have positive duration. Since it is countable, the set of all positive-duration members of Γ has only countably many finite subsets. So, the equivalence relation that holds between two instants when they instantiate exactly the same positive-duration properties in Γ partitions the set of all instants into countably many cells. Assuming there are more than one of these cells, we then will have to say something about when the transitions between cells occur. Given that the changes in the relevant underlying facts are gradual, one might worry that there is no principled explanation available of why the transitions should take place when they do, rather than at any other

time. An analogous worry can be raised in the modal case: if the space of physically accessible worlds divides into countably many cells, such that within each cell exactly the same positive-measure propositions from a given family Γ are true, anything we might say about the locations of the cell-boundaries will seem objectionably arbitrary.

We definitely feel the force of these worries. However, turning them into a compelling argument would be a delicate task, since there is a risk that the required premise will on close inspection turn out to be tarred with the brush of the Sorites. After all, we who accept theorems of classical logic must believe that there is a countable partition of the time-series such that t and t' belong to the same cell if and only if the same people are bald at t and at t' , even though, in view of the continuity of the relevant physical processes, any particular proposal about when the transitions between the cells of this partition occur will likewise seem terribly arbitrary.²³ By contrast, the arbitrariness worry that drives the case for *Parity* is fairly specific to the domain of semantic facts (and perhaps also psychological facts). For the propositions expressed by ordinary sentences not about semantics—say, the proposition that there are dogs, or that salad is delicious—there is no obvious way of specifying an uncountable family of similar propositions for which there would be any temptation at all to think that only finitely many members of the family are true at typical nearby worlds.²⁴

²³ Even if considerations of arbitrariness do not provide good reason to deny that there is a least number of hairs consistent with non-baldness, they may provide good reason to deny that any number is *definitely* the least number of hairs consistent with non-baldness, and may similarly provide good reason to deny that any pair of close times or worlds is *definitely* separated by a cell-boundary. §4.4 will develop a new argument from *Finitude* to *Parity* based on this claim.

²⁴ It may be worth mentioning a third kind of arbitrariness-based worry, distinct both from the one that supports *Parity* and from the one that drives arbitrary-boundaries arguments. Given the supervenience of the semantic on the microphysical, there must be some function that characterizes the way semantic facts are fixed by microphysical facts. Why—one might ask—should *that* function, rather than some other very similar function, be the one that specifies the semantic facts? To the extent that there is anything to this concern, it is extremely general. It has nothing in particular to do with

3. Embracing plasticity

3.1 *Can modal plasticity be reconciled with counterfactual speech reports?*

Suppose that we conclude that the propositions expressed by typical speech reports have zero phase-space measure. Then we will face a pressing challenge when it comes to counterfactuals like (2):

(2) If we had taken Sally to Giorgio's, she would have said that salad is delicious.

We see no reason to deny that the antecedent of (2) has positive phase-space measure.

So if its consequent has zero phase-space measure absolutely, it also has zero phase-space measure relative to its antecedent. This makes (2) disturbingly analogous to utterly unassertable counterfactuals like

(5) If I had thrown a dart at this dartboard, it would have hit the exact centre.

After all, the community of English-speakers would have to have been different in many microphysical respects for us to have taken Sally to Giorgio's. If the proposition that Sally asserts that salad is delicious is so ultra-demanding that it is true only in a zero-measure set of possible microphysical configurations, how could we ever reasonably be confident that one of these very special configurations would have obtained if we had taken Sally to Giorgio's?²⁵

semantic plasticity, or indeed with semantics at all: it applies to any subject matter where there seems to be supervenience in the absence of short, natural, necessary and sufficient conditions. Unless one is prepared to give up supervenience for all "messy" subject matters, one will need somehow to make one's peace with arbitrariness worries of this third sort.

²⁵ The basic worry here also arises for the view that the proposition that Sally said that salad is delicious has a positive, but extremely small, phase-space measure relative to the proposition that we take Sally to Giorgio's: in that case, the disturbing analogy is with counterfactuals like 'If I had thrown a dart at this dartboard, it would have come within such-and-such microscopically small distance of hitting the centre'. However, some of the possible responses to the worry are more promising if the measures are small but positive than if they are zero: see footnote 29 below.

In pressing this challenge, it would be a mistake to draw from (5) the overly general lesson that every counterfactual whose antecedent is known to be false and whose consequent has zero measure relative to its antecedent is unassertable. This principle is false. For example, suppose certain particles *a* and *b* are too distant from Sally for any kind of causal relation to obtain between her and them.²⁶ Then (6) seems acceptable, despite having a positive-measure antecedent and a zero-measure consequent:

- (6) If we had taken Sally to Giorgio's, *a* and *b* would have been exactly as far apart as they actually are.

On the most influential account of counterfactuals (Stalnaker 1969, Lewis 1973), our confidence in (6) could be vindicated by appealing to our right to be confident that among worlds where we take Sally to Giorgio's, at least one where *a* and *b* are exactly as far apart as they actually is more similar to the actual world (according to the contextually appropriate metric of similarity) than any where they are a different distance apart. An alternative way to vindicate counterfactuals like (6) appeals to the idea that on any given occasion when a counterfactual is uttered, the domain of relevant worlds is restricted to worlds which match the actual world in certain contextually given respects (von Fintel 2001). To explain the acceptability of (6), we could simply say that in the context of a typical utterance of (6), the truth about the distant region of the universe containing *a* and *b* is held fixed, so that the only worlds that need to be considered are those which match actuality as regards the state of this region.²⁷

²⁶ In a relativistic setting, this assumption could be precisified as the claim that *a* and *b*'s past light cones do not overlap Sally's.

²⁷ In this case, the domain of worlds that match actuality in the relevant respects is itself a zero-measure set. But this is not required for the holding-fixed approach to save a counterfactual whose consequent has zero measure relative to its antecedent. Consider the following example, which we owe to Timothy Williamson: 'If the dart had either hit the exact centre of the dartboard or missed altogether, it would have hit the exact centre'. This is arguably acceptable when the dart has in fact hit near the centre, although the case is less clear than for (6). We could account for this by claiming that the domain of

It seems to us that one or other of these ideas provide the best hope for saving the view that (2) is assertable despite having a zero-measure consequent. In the similarity-theoretic setting, a flat-footed way to save (2) would be to build certain semantic facts *de jure* into our account of the relevant similarity measure. This means that worlds that match actuality with regard to certain zero-measure semantic facts, such as the fact that 'salad' means *salad*, will thereby receive an automatic boost in the similarity ranking vis-à-vis other worlds that may be just as close to the actual world as them according to any natural microphysical measure. An alternative, more ambitious approach would be to characterise the operative similarity metric in non-semantic terms, and then argue that in very many cases there is reason to expect the worlds that do best by this standard to belong the relevant zero-measure sets. We can likewise distinguish two ways of pursuing the holding-fixed strategy. On the one hand, we could flat-footedly incorporate certain zero-measure semantic facts, such as the fact that 'salad' means *salad*, into the contextually specified body of facts that are held fixed, so that worlds which do not match actuality in those semantic respects are automatically excluded as irrelevant. On the other hand, we could ambitiously attempt to *explain* the exclusion of those worlds by postulating that certain non-semantic facts about actuality are held fixed, and arguing that despite having absolute measure zero, the consequents of the problematic counterfactuals have high measure relative to the sets of worlds where those non-semantic facts obtain and the antecedents are true.

relevant worlds contains all and only worlds that match actuality with respect to whether the dart hits the board. This set has positive-measure, although its intersection with the set of worlds where the antecedent is true still has zero measure. There is also a natural similarity-theoretic account of the acceptability of Williamson's counterfactual, in which the relevant similarities have to do with where the dart lands; however, it is a challenging task to work this out without being forced to endorse bizarre counterfactuals like 'If the dart had landed two centimetres or more from where it actually landed, it would have landed exactly two centimetres from where it actually landed'.

For now, let us make the plausible assumption that in evaluating counterfactuals, we do not venture outside the set of physically accessible worlds unless the antecedent forces us to: provided that the antecedent is consistent with the actual laws, the actual number of particles, etc., the worlds relevant to the consequent also match actuality in these respects. (We will consider dropping this assumption in §3.2.) Given this assumption, the prospects for the ambitious versions of the similarity-theoretic strategy and the holding-fixed strategy seem dim. Let us consider the holding-fixed strategy first. For the holding-fixed strategy to have any chance, at least one zero-measure fact must be held fixed.²⁸ What could this be, if it is not simply a semantic fact? It is plausible that, in evaluating a counterfactual about goings on at a certain time t , we hold fixed many truths about history before t . This helps explain the acceptability of counterfactuals like

- (7) If John had forgotten to have breakfast this morning, that would have been the first time he did so in months.

But remember that, given our current assumptions, we cannot hold fixed the historical truths in every microphysical detail: since the laws are deterministic and must be held fixed whenever possible, we have to say that if we had taken Sally to Giorgio's, history would have been at least *slightly* different right from the beginning. The most we can do is to hold fixed some restricted collection of historical truths. For example, on one approach, we hold fixed the *macroscopic* truths about history up to the relevant time. This is perfectly adequate for explaining the acceptability of ordinary counterfactuals like (7). But this is of no use for present purposes, since in the sense of 'macroscopic truth' that figures in statistical mechanics, the conjunction of all macroscopic truths will

²⁸ In principle the conjunction of the antecedent with the held-fixed facts could have zero measure even if neither the antecedent nor the conjunction of the held-fixed facts has zero measure by itself (see the previous footnote). But this seems like a lot to hope for, given that the antecedents of the counterfactuals we are interested in are not disjunctive in the way that the proposition that the dart either hits the exact centre of the dartboard or misses altogether is disjunctive.

always have positive measure.²⁹ It would not be unreasonable to suggest that certain zero-measure non-semantic truths are held fixed in addition to these positive-measure truths: for example, the truth about the total energy of the universe, or about the states of extremely distant portions of the universe. This would explain the assertability of counterfactuals like (6). But these 'physically natural' zero-measure truths are also not going to help save counterfactuals like (2). For as we saw in §2, the considerations that support the conclusion that the set of physically accessible worlds where Sally says that salad is delicious has zero measure *simpliciter* also support the conclusion that it has zero measure relative to other 'physically natural' zero-measure sets of worlds, such as the actual energy hypersurface. It is thus completely obscure what facts, other than the semantic facts themselves, could play the relevant explanatory role.

For analogous reasons, we think that the ambitious version of the closeness strategy is unlikely to succeed. Seen through the lens of physics, the zero-measure sets of worlds

²⁹ This is one place where it matters that we are focusing on the very strong claim that typical speech reports have measure *zero*. If one rejected this in favour of the claim that such reports have very small positive measure (perhaps because of the physical and metaphysical worries about continua discussed in footnote 20), there would be more room to hope that the holding fixed of macroscopic truths about history would vindicate counterfactual speech reports like (2). For since the complete macroscopic truth about history also has very small measure, one could consistently hold that the measure of the proposition that Sally asserts that salad is delicious is very small absolutely but large relative to the set of worlds where we take Sally to Giorgio's and the macroscopic truth about earlier history is just as it actually is. However, we are not very optimistic about this as a way out of our puzzle. To work, it would need the kind of "backwards-looking" metasemantics we will sketch in 3.2: a metasemantics on which words change their meanings quite infrequently, so that worlds which match macroscopically up to *t* will typically agree about the meaning of any given word for quite a long time after *t*. But for many reasonable-seeming counterfactual speech reports, the time of the reported speech is long after the time at which the relevant worlds must diverge macroscopically from actuality: consider 'If you had made me eat my vegetables as a child, I would be telling you that that was a delicious salad', or 'If Caesar had had a limp, twentieth-century history textbooks would have mentioned the fact'. It is hard to see how, without giving facts about the distant past an absurd level of metasemantic importance, one could plausibly hold that the measures of the consequents of these counterfactuals are very small absolutely but not small relative to the sets of worlds matching the actual world macroscopically until the times mentioned in the antecedents.

where semantic propositions are true will plausibly have quite intricate, arbitrary-looking shapes in phase space. Given this, it would be extremely surprising if we could, without bringing in semantic considerations, characterise in physical terms a closeness measure on which worlds matching actuality in certain zero-measure semantic respects reliably counted as closer than neighbouring worlds not matching the actual world in those respects.

This leaves us with the flat-footed versions of the closeness and holding-fixed strategies, in which certain semantic facts are built *de jure* either into the similarity metric or into the body of facts held fixed. Implementing either of these strategies will be a delicate task, since there is a serious risk of overgeneration. Suppose that unbeknownst to us, certain malevolent aliens have taken an interest in salad-loving Sally, and decided that if she ever enters Giorgio's, they will activate a certain fairly reliable device which functions to completely rewire everyone's linguistic dispositions. If we get to hold fixed the fact that 'salad' means *salad* in evaluating our utterance of (2), we will get the undesirable result that this utterance expresses a truth in spite of the aliens, since the only relevant worlds where the antecedent is true are ones where the aliens' device happens not to work. In principle this problem could be avoided by finding some weaker (but still zero-measure) proposition to hold fixed, but it is unclear what this would be. Likewise, it is far from straightforward to construct a measure of similarity that gives a role to semantic considerations without making worlds where the aliens' device misfires count as more similar to actuality than worlds where it works, and thereby counting (2) as true even in the world with the aliens.

Another pressing concern about the flat-footed strategies is that, in looking beyond considerations of physics in deciding what is held fixed or what makes for similarity, they are liable to override probabilistic considerations derived from physics in

objectionable ways. For example, suppose that Fred is both an excellent salad-maker and one of the least cautious people around when it comes to the use of the word 'salad': he unreflectively calls dishes 'salads' even in cases when most others would be reluctant to apply the word. One year ago, Fred's enemy Zack was planning to shoot Fred, but had a change of heart at the last minute. Fred went on to exercise his liberalism about 'salad' on many more occasions. Consider the following counterfactual:

- (8) If Zack had shot at Fred, he would have killed him, and right around now we would be having a memorial service at which people would be saying how delicious Fred's salads were.

Given appropriate background assumptions—that Zack is a good shot; that his gun is very reliable; that Fred's salads are indeed so memorable that people would be apt to dwell on them at a memorial service...—(8) could be a reasonable thing to assert. Of course, Zack is not *infallible*—for example, Fred's life could have been saved if the bullet had been blown off course by an unlikely gust of wind. But (*pace* Hájek MS) a counterfactual can be acceptable even when the antecedent does not raise the chance of the consequent all the way to 1. However, if securing the truth of the speech report in the consequent of (8) requires holding fixed certain zero-measure facts about the current meaning of 'salad', or invoking a similarity measure that gives a boost to worlds where these facts obtain, then our right to be confident in (8) may come under considerable pressure. For while the physical facts about Zack and his gun make it plausible that the proposition that Zack kills Fred has high measure relative to the set of worlds where Zack shoots and earlier macro-history is the same, this proposition may well have a much lower measure relative to the set of worlds where Zack shoots, earlier macro-

history is the same, and 'salad' still ends up meaning *salad*.³⁰ For 'salad' to continue to mean *salad* despite Fred's absence would require some kind of compensating adjustments in the linguistic behaviour of other speakers. And there is no obvious reason to expect that worlds where these adjustments occur have higher measure, within the relevant set, than worlds where Zack misses and Fred survives. Remember that the zero-measure region of phase-space corresponding to worlds where 'salad' means *salad* will be quite strangely shaped, with all kinds of oscillations and curlicues that look quite unnatural from a physical point of view. It would thus be quite unsurprising if the measures of many propositions relative to this zero-measure region were significantly different from their measures relative to any relevant region picked out in a physically natural way. To make a reasonable estimate of the measure of the consequent of (8) relative to this region, one would need to take into account many facts about the physical processes in the rest of the population that might have led to the required compensating adjustments, in addition to the facts about the physical workings of Zack's gun. And the point is quite general: if in evaluating certain counterfactuals we have to hold fixed zero-measure semantic truths or appeal to a similarity measure that

³⁰ A fully rigorous reconstruction of our argument in this paragraph would involve a premise connecting the phase-space measure with credences in ordinary counterfactuals. The following premise would do the job: for a high confidence in 'If it were that P it would be that Q' to be appropriate, it must be reasonable for you be confident that the proposition that Q has a high measure relative to the set of *relevant* P-worlds, where the relevant P-worlds are those which match actuality in the contextually relevant respects and/or are maximally close to the actual world by the operative similarity measure. This premise is a close cousin of the 'Ramsey Bound' principle discussed in Williams 2012. Williams develops an argument, modelled on one of the triviality proofs in Lewis 1976, that the Ramsey Bound must be rejected if counterfactuals have truth conditions at all. We think Williams's argument can be resisted by being careful about the context-sensitivity of counterfactuals, but this is not the place for a full rebuttal.

We are not sure that the phase-space measure really does play the role in governing the acceptability of counterfactuals that the above premise attributes to it. But the argument will work just as well so long as some physically natural measure plays this governing role—for example, the relevant measure could be one derivable from the phase space measure by some operation that concentrates probability on the set of worlds where some 'low entropy past hypothesis' is true (Albert 2000).

gives them special weight, then our ordinary practices for evaluating those counterfactuals will be apt to lead us astray, since these practices take no account of the probabilistic dependencies between those semantic propositions and the non-semantic propositions that occur in the counterfactuals' antecedents and consequents.³¹ While this conclusion is not as disastrous as the claim that we should assign zero credence to most ordinary counterfactuals whose consequents contain speech reports, it is nevertheless quite unattractive—enough to warrant a close look at other approaches to the puzzle.

3.2 *Miracles to the rescue?*

§3.1 assumed that ordinary counterfactuals with nomologically possible antecedents do not take us to nomologically impossible worlds, even if determinism is true. But there is a well-known alternative to this picture. According to Lewis (1979), if determinism is true, then if John had forgotten to have breakfast this morning, times earlier than this morning would have been the same, not just macroscopically but down to the last microphysical detail. Then there would have been a small 'miracle'—an isolated counterexample to the actual laws of nature—after which the further exceptionless playing out of the actual laws would have involved John's forgetting breakfast.³² This

³¹ Jeremy Goodman (in conversation) and an anonymous referee both suggested getting around this problem by holding fixed the truths not just about semantics, but about some even finer-grained subject matter, cunningly chosen in such a way as to guarantee that the physical considerations we normally take into account *are* a good guide to the measures of the propositions we are interested in relative to the (zero-measure) set of worlds that match actuality with respect to that subject matter. There is no obvious way of specifying a subject matter with this convenient feature. But perhaps, as Goodman suggested, the theory of vagueness could come to rescue here: it may be vague what the subject matter is, but definite that it has the profile required for our usual methods to be a good guide to figuring out which propositions have high measure relative to the set of worlds that match actuality with respect to it. (This will require some account of vagueness on which one can reasonably have a high credence both that P and that it is vague whether P.) We have no argument against this intriguing proposal, although it strikes us as far-fetched.

³² A central motivation that Lewis gives for this view is the worry that, given determinism, there might be *no* nomologically possible world which almost matches the

view makes *Plasticity*, which concerns how the semantic facts work in nomologically possible worlds, irrelevant to the truth-values of counterfactuals like (2)–(4). It thereby opens up the possibility of a metasemantics that allows such counterfactuals to be true without the need for semantic facts to be *de jure* held fixed.

Here is one toy model of such a metasemantics, which combines modal plasticity with temporal non-plasticity. Words change their meanings in sudden jumps. After each jump, use-based pressure for the next jump starts to build up, but the meaning remains constant until this pressure reaches a certain critical threshold, which typically takes several decades. Which meaning a word ends up with after a jump depends only on previous history: this is a “backward-looking” metasemantics.³³ Suppose that the last jump in the meaning of ‘salad’ was five years ago. Then even if Zack had killed Fred last year, the effects of Fred’s absence would not have manifested themselves semantically for many more years. Thus insofar as we can be confident that we are not

actual world up to a given time, and then diverges in such a way as to make true the antecedent of a counterfactual concerning that time. Considering the hypothesis that this is how the world would have been if Nixon had pressed a certain button, Lewis claims that ‘there is no guarantee whatever that [a nomologically possible world where Nixon presses the button] can be chosen so that the differences diminish and eventually become negligible in the more and more remote past. Indeed, it is hard to imagine how two deterministic worlds anything like ours could possibly remain just a little bit different for very long. There are altogether too many opportunities for little differences to give rise to bigger differences.’ (Lewis 1979: 45) But the worry Lewis expresses in the last two sentences is misplaced. Our best deterministic physical theories have continuous dynamics: for each Δt , the function that takes a given point in phase-space and evolves it forwards by Δt units of time is continuous. Supposing there is a first moment of time, this means that for any n and ϵ , we can find a perturbation of the state at the first moment that is so small that even after n years, the state of the system will still be within distance ϵ of its actual position (according to any well-behaved way of measuring distance in phase space). Of course, there are other worries that could be raised about the viability of an approach to counterfactuals that holds the laws exactly fixed and the past approximately fixed, but this is not the place for a full defence of such an approach. (For some recent accounts of counterfactuals under determinism in this spirit see Albert 2000 and Loewer 2007.)

³³ On the ‘temporal externalist’ view that physical facts about later times make a difference to semantic facts at earlier times (Jackman 1999), it is hard to see how a Lewisian approach to counterfactuals could provide any help with our puzzle.

too near a jump, we can be confident that people nowadays would still have meant the same thing by 'salad' if Fred had been killed. In general, the model allows us to claim that homophonic counterfactual speech reports whose their consequents are about times not too much later than those their antecedents are about are mostly true. (Counterfactuals like 'If Caesar had had a limp, twentieth-century history textbooks would have mentioned the fact' are a lot harder to save.)³⁴

No theory of how counterfactuals work under determinism is pain-free; Lewis's view is one of the most worked out approaches, and its apparently superior capacity to save counterfactuals with semantic consequents constitutes a novel and substantial count in its favour. Nevertheless, we think the view's strangeness has been underestimated. Suppose that some simple deterministic system of laws *L* is in fact correct, and that Frank, a philosopher of physics, has devoted his career to defending the (true) claim that the laws in *L* hold universally and without any exceptions whatsoever. Frank has a debate with an opponent, Nancy, who maintains (wrongly) that there are exceptions to *L*. On Lewis's view, if the circumstances of the debate had been different in any way whatsoever—for example, if an extra glass of water had been put on the lectern—Frank would have been wrong, and Nancy right. We find this really hard to live with. If Frank knows how counterfactuals work under determinism, should he be saying to himself 'Thank goodness they didn't put an extra glass of water out, otherwise my whole career would have been based on a mistake?'

Of course, our preferred view, according to which facts all the way back to the Big Bang would have been microscopically different if an extra glass of water had been on the lectern, also takes some getting used to. But any microscopic truth about the distant

³⁴ One could also appeal to a metasemantics along these lines to explain how homophonic counterfactual speech reports get to be true in indeterministic worlds, even if one rejects Lewis's account of counterfactuals under determinism.

past that nomically necessitated some truth about the present distribution of glasses of water would be far too complex for anyone ever to express in physical terms. So there is no straightforward way to construct an analogue of the above argument in which Frank's career has been devoted to defending one of the truths about the distant past which are counterfactually fragile according to us but not according to Lewis.

3.3 *Can contextual inheritance rescue homophonic reports?*

The strategies for blocking the objection to modal plasticity from counterfactual reporting which we considered in §3.1 and §3.2 turned on certain controversial claims about the semantics of counterfactuals. Neither of these strategies suggests any way to block the objection to *temporal* plasticity from cross-temporal reporting. In this subsection and the next, we will consider strategies which do not depend so crucially on the details of the semantics for counterfactuals, and which are potentially of use in defending temporal plasticity as well as modal plasticity.

We have been writing as if there were such a thing as *the* proposition expressed (in English, as spoken now at the actual world) by 'Salad is delicious'. But many sentences are not like that: at any time, a *context-dependent* sentence is semantically associated with a range of propositions, different members of which are asserted or expressed on different occasions when the sentence is used. It would not be absurd to think that the considerations that make for semantic plasticity always also make for context-sensitivity. For those who think that this is the case, a new strategy opens up for vindicating cross-modal (and cross-temporal) speech reports while accepting that their consequents are subject to modal (and perhaps also temporal) plasticity.

Let us take 'local' as our paradigm context-sensitive word. Familiarly, its semantic contribution on a given occasion is not mechanically anchored to the locale of utterance. When other people or places are salient, 'local' can contribute the property of being

local to them. And this is often the case when 'local' appears in speech reports. For example, 'Fred said he was going to a local bar' is naturally taken to mean that Fred said he was going to a bar local to *him*. If 'salad' had an analogous flexibility, then we could salvage (1) ('Sally said that salad is delicious') even in the face of widespread temporal plasticity by taking this particular occurrence of 'salad' to inherit the content that 'salad' contributed to Sally's utterance. Even if two speakers neither of whom is salient to the other rarely or never express the same content with 'salad', it could be that when a later user of 'salad' is focusing on a particular earlier use of 'salad' (or some other word with a similar semantic profile), the later use inherits the content of the earlier use.³⁵ Similarly, we might be able to salvage counterfactual reports like (2)–(4) by holding that the context-sensitive words in the report have their contents fixed in such a way as to match the content expressed by the target in the closest world where the antecedent of the counterfactual is true.

This way reconciling plasticity with cross-temporal and counterfactual reports has sharp limitations. It does not work for *quantified* cross-temporal reports like

(9) All new customers who came to Giorgio's said that their salad was delicious

assuming that the target speakers were not all focusing on some other utterance in the way that makes for content-inheritance. Similarly, inheritance is insufficient to account for counterfactual reports where a range of speakers and/or worlds are involved. Indeed, if we follow Lewis (1973) instead of Stalnaker (1968) by not requiring a unique closest world, even simple counterfactuals like (2)–(4) will fall into this category. And by anyone's lights, (10) will be like this:

³⁵ It is important that the mechanism of "inheritance" doesn't only occur when context-sensitive items are embedded in attitude reports. Imposing that restriction would disrupt such inferences as 'She correctly said that salad is delicious; therefore, salad is delicious.'

- (10) No matter which restaurant we had gone to, Sally would have asserted that salad is delicious.

The problems speech reports involving multiple times and worlds pose for an defence of plasticity based on contextual inheritance are similar to the problems certain quantified sentence involving 'local' pose for the view that each occurrence of 'local' refers to some particular locality. For example, one can speak truly by saying 'Each of my brothers is going to a local bar tonight' even if one's brothers are scattered around the globe. There is no locale such that anchoring 'local' to it will generate a true reading for this sentence. In this context, 'local' behaves as if it included a variable bound by a higher quantifier (in this case, 'each of my brothers').³⁶

To reconcile quantified reports like (9) and (10) with temporal and modal plasticity, one could propose that 'salad' works like 'local' in this respect. The occurrences of 'salad' in (9) and (10) behave as if they contained hidden variables bound by external quantifiers. When one utters (10), there is no particular proposition p such that one asserts that no matter which restaurant we had gone to, Sally would have asserted p : rather, one asserts that no matter which restaurant we had gone to, Sally would have asserted whichever proposition she would in that case have been mapped to by a certain function. Provided that we are confident that Sally would have uttered the sentence 'Salad is delicious' no matter which restaurant we had gone to, and that the relevant function is one that maps people uttering 'Salad is delicious' in nearby worlds onto propositions that they thereby assert, our right to be confident in (10) will be secured.³⁷

³⁶ See Partee 1989, Cappelen and Hawthorne 2009, ch. 2.

³⁷ Note that the proposition that Sally asserts whichever proposition she is mapped to by the relevant function—which is the proposition that 'Sally asserts that salad is delicious' contributes in the context of (10)—can be positive-measure even if, for each proposition P , the proposition that Sally asserts P is zero-measure.

Even if these techniques allow contextualists to generate readings for the problem sentences on which we can be confident that they are true even if temporal and modal plasticity is widespread, we might reasonably object that the semantics also predicts the availability of *false* readings of the problem sentences, and thus true readings of their negations, which we seem unable to access. By contrast, we have no trouble hearing the reading of 'None of my brothers will be in any local bar tonight' on which it is true if none of the speaker's brothers will be in any bar local to the speaker. Normally, when a sentence is structurally ambiguous, and comes out true on one way of resolving its structural ambiguity, considerations of charity will often lead us to interpret people who utter the sentence as intending this resolution. In the case of 'local', we find exactly this: the 'sloppy' readings where it works as if it contained a bound variable and the 'strict' readings anchored to a particular locale are readily distinguished, and we choose between them using considerations of charity and contextual cues. Moreover, there are various linguistic environments in which the 'strict' reading is overwhelmingly natural. For example it is hard, and perhaps impossible, to access a reading for 'There is a certain bar, namely the local bar, that each of my brothers will be in tonight' on which it can be true when the brothers will be in different cities. These facts militate against a treatment of cross-temporal and cross-world reporting that depends on bound-variable readings for worlds like 'salad'. Hearing (2)–(4), we have no sense whatsoever of a relevant structural ambiguity that needs to be resolved. If someone utters

(11) No-one has ever said that salad is delicious

we do not notice an option of saving the truth of the speech by adopting a "strict" reading on which nothing in 'salad is delicious' is bound by 'no-one'. And even if one can somehow generate a "sloppy" reading for (12),

- (12) There is a certain claim, namely that salad is delicious, which Sally feels so strongly about that she would certainly have propounded it no matter which restaurant we had taken her to,

it strikes us as highly implausible that ordinary speakers either invariably fail to notice the false, binding-free reading, or that they notice it but never intend it. (The argument that (9) is false on this reading is by no means so straightforward that the folk are likely to have already been swayed by it.)

3.4 *Loose talk?*

The prospects for reconciling pervasive modal plasticity with the truth of ordinary counterfactual speech act reports do not look good, and the prospects for reconciling pervasive temporal plasticity with the truth of ordinary diachronic speech act reports look even worse. In §1, we suggested it would be a 'disaster' to have to reckon typical reports in either category as false. Is this right, or can the blow be softened to a tolerable level?

An important choice point confronts any view according to which some wide range of ordinary utterances are false. Do ordinary people typically *believe* the false propositions they semantically express in making those utterances? Or are they taking liberties, making utterances that express propositions that they do not believe? Call the former approach *error-theoretic* and the latter approach *pragmatic*.

We have little to say to error-theoretic bullet-biters, except that we, like most analytic philosophers, tend to be alarmed when philosophers claim to have found arguments that overthrow large swathes of ordinary peoples' belief system. (Those who do make such claims face familiar *ad hominem* charges when, outside the philosophy class, they revert to the very ways of talking whose prevalence among non-philosophers is taken as a sign of false belief.)

The central challenge facing proponents of a pragmatic approach is to tell a systematic story about how the choice to utter one sentence rather than another constrains what one can get across by an utterance, given that one is not trying to get across the proposition semantically expressed. A generic way to respond to this challenge is to assign the problematic sentences “loose contents” distinct from the propositions they semantically express, such that when we assert a sentence in an ordinary setting, we must really believe the proposition that is its loose content unless we are being dishonest.³⁸ What, then, does it take for the loose content of (2) to be true? It seems very plausible that it has *something* to do with how things would have been if we had taken Sally to Giorgio’s. The divergence between strict and loose contents is not a matter of our caring about some *different* counterfactual scenario or scenarios other than the one where we take Sally to Giorgio’s: rather, it is due to there being some other proposition, distinct from the proposition that Sally says that salad is delicious, such that in asserting (2) we are really committed to believing that *it* would have been true if we had taken Sally to Giorgio’s. Let us use the sentence ‘Sally quasi-asserts that salad is delicious’ to semantically express the proposition that plays this role. More generally, let us use ‘quasi-assert’ in such a way that, when certain sentences Φ and Ψ do not themselves involve any semantic vocabulary or other expressions that give rise to a divergence between loose and strict content, the proposition semantically expressed by ‘If it had been that Φ , α would have quasi-asserted that Ψ ’ is equivalent to the loose

³⁸ It would be a mistake to assume that the loose content is always weaker than the strict content. If ‘The table is flat’ is strictly false and loosely true, ‘The table is not flat’ is presumably strictly true and loosely false. Presumably the notion of loose content will need to be relativised to context in the same way that the notion of semantic expression normally is; but for simplicity’s sake we will suppress all the relevant relativizations in the present discussion.

content of 'If it had been that Φ , α would have asserted that Ψ '.³⁹ (On one natural proposal, to quasi-assert a proposition is just to *assert* a proposition that is "close enough" to it in certain respects. But a theory of loose content certainly need not take this form for our introduction of 'quasi-assert' to be successful.)⁴⁰

One major problem faced by any philosopher who appeals to loose content is to explain why loose contents do not, after all, count as being semantically expressed. The risk is that the concept of semantic expression will turn into a mere theoretical idler, which does not link up with other concepts related to communication in such a way as to give it any explanatory point. This risk will be more serious to the extent that the assignment of loose contents to sentences is determined by their syntactic structures according to systematic rules, since the idea that semantic contents are determined in this way is a crucial part of semantic content's explanatory role. Take for example the view on which the loose content of (2) is true if and only if, had we taken Sally to Giorgio's, she would have asserted something close enough to the proposition that salad is delicious: the challenge for proponents of this view is to explain why ' x asserts p ' has the semantic content they assign to it rather than the one they assign to ' x asserts something close enough to p '. But even leaving such general issues aside, it is important to see that the pragmatic approach does not allow one to completely sidestep the challenge posed by the continuum argument. Having introduced expressions like 'quasi-assert', we can construct new instances of the continuum argument, where the target

³⁹ We can do the same thing for other speech-act verbs, so we can speak of 'quasi-saying', 'quasi-promising', 'quasi-conveying', etc.

⁴⁰ Our introduction of 'quasi-assert' does not explicitly stipulate anything about its semantic content, or the semantic content of sentences involving it other than 'If it had been that Φ , α would have quasi-asserted that Ψ '. Our intention is that 'quasi-assert' should make whatever semantic contribution it needs to make for the semantic contents of those sentences to be determined compositionally. This may not uniquely characterise a single meaning; if it does not, 'quasi-assert' will have the kind of vagueness characteristic of predicates introduced by stipulations that fall short of being complete definitions. This should not matter for our purposes.

family of propositions Γ concern quasi-assertion rather than assertion. And accepting these new instances of the continuum argument will not be an option for proponents of the loose-talk approach, given that what motivates their view is their acceptance of the inference from the claim that certain propositions are subject to modal plasticity to the conclusion that we cannot reasonably be confident that those propositions would have been true under various ordinary counterfactual circumstances. So, one or other premise of the continuum argument for quasi-assertion must be false. Either *Parity* fails, so that there are big differences in the ease with which certain propositions can be quasi-asserted in spite of their physical and metaphysical similarities, or else *Finitude* fails, so that it is possible for us to quasi-assert infinitely many propositions in a single utterance. (Those who equate quasi-asserting a proposition with asserting something close enough to it will probably reject *Finitude*, since in typical cases there will be infinitely many propositions that count as 'close enough' to a given proposition.) As we will see in the remainder of the paper, those who respond to the original continuum argument by denying *Parity* or *Finitude* face a multitude of challenges, and for most of these, a switch from assertion to quasi-assertion will do nothing to smooth things out.

4. Rejecting *Parity*: the patchwork model

As we have seen, the conclusion that everyday speech reports express zero-measure propositions is quite hard to maintain. And the temporal analogue of that conclusion—according to which the act-types described by everyday speech reports typically have at most one instance in the history of the world—has a weaker motivation and is even harder to maintain (since the proposals about the semantics of counterfactuals discussed in §3.1 and §3.2 lack temporal analogues). We are thus left with a choice between rejecting *Parity* or rejecting *Finitude* (or both). In this section and the next, we will

investigate these options. As we will see, even if we are willing to set aside §2's arguments for *Parity* or *Finitude*, it is far from straightforward to see how either of these strategies can provide an adequate vindication of our cross-temporal and counterfactual reporting practices.

The present section will focus on what we will call 'the patchwork model', which combines the denial of *Parity* with the following orthodox view about assertion, which provides the most obvious motivation for *Finitude*: in a normal assertoric utterance, exactly one proposition is asserted, namely the unique proposition semantically expressed by the sentence uttered relative to the context of the utterance. Many of the problems we will discuss will also apply to *Parity*-denying views on which typical acts of assertion involve the assertion of more than one proposition, but not infinitely many; but we will postpone discussion of such views until §5, since the reporting-related challenges they face are similar to those facing deniers of *Finitude*.

4.1 *Limited errors in counterfactual and cross-temporal reports*

So long they are willing to make the patches of semantic homogeneity reasonably large, proponents of the patchwork model will have no trouble reckoning *most* homophonic reports of close temporal and modal neighbours as true. But some such reports, namely those where the relevant sentence expresses one proposition at the time and world of the original utterance and expresses a different proposition at the time and world of the report, are inevitably going to have to be counted as false. For example, let t_1 be some time shortly before 'Salad is delicious' came to mean that salad is delicious, and let t_2 be some time shortly after the shift. Suppose Sally uttered 'Salad is delicious' at t_1 , prompting Fred to utter 'Sally just asserted that salad is delicious' at t_2 . Since 'Salad is delicious' meant something different in Sally's mouth from what it meant in Fred's (and means in ours), it seems to follow that Fred's report is false. And since Fred was

speaking literally and sincerely, the false proposition he asserted is one that he *believes*. Although the method he used to form this belief is normally reliable, in this case it led him into error.⁴¹

The idea that the method of homophonic reporting sometimes leads us astray in cases like this one takes a lot of getting used to. Of course, no-one would want to say that the method is infallible. Sometimes, we wrongly assume that someone who uttered a certain sentence Φ used it to assert the same thing that *we* would have asserted if we had uttered Φ , and we express this false belief when we say 'So-and-so asserted/claimed/said that Φ '. However, the cases where it is intuitive to posit such errors are generally ones where the reporters wrongly take themselves to be more similar to the subjects than they in fact are. Fred's putative error is not like that: Sally at t_1 really is *extremely* similar to Fred at t_2 in all the ways that matter, and Fred is under no misapprehension about this: it's just that the tiny difference between Sally at t_1 and Fred at t_2 happens, unluckily for Fred, to be one of the few such differences that makes for a semantic difference. It is hard to believe that we are in danger of going wrong in the way Fred goes wrong about what is said by people whom we know to be so similar to ourselves.

The argument given above for the conclusion that Fred said something false contains an important flaw. To derive that conclusion from the premise that Sally did not assert that salad is delicious, we need to assume that in uttering 'Sally asserted that salad is delicious', Fred asserted that Sally asserted that salad is delicious. Given the orthodox view of assertion, this will be true so long as, at t_2 , Fred's sentence *meant* that

⁴¹ One resource for blocking this result is one that we have already explored on behalf of *Plasticity*, namely to posit context-sensitivity in words like 'salad', together with a mechanism of inheritance whereby Sally's salience to Fred makes Fred's use of 'salad' express the same content as Sally's. But for reasons explained in §3.3, we doubt that this contextualist strategy is adequate to save all the reports that need to be saved. If we are wrong about this, there is little motivation for denying *Plasticity* in the first place.

Sally asserted that salad is delicious. But this does not follow from our stipulation that at t_2 'Salad is delicious' meant that salad is delicious, since that stipulation is consistent with the hypothesis that 'assert' has shifted its meaning between t_2 and now. Perhaps Fred was not talking about assertion, but about some other relation, *assertion**. And perhaps, although Sally didn't assert the proposition that salad is delicious, she did *assert** it, so that what Fred asserted is true. Presumably, given that *assertion** is similar to assertion, there is a relation—call it *meaning**—that stands to *assertion** as *meaning* stands to assertion: normal assertoric utterances involve the *assertion** of exactly one proposition, namely the unique *meaning** of the sentence uttered. The meaning relation determines one partition of the space of times and worlds, in which points belong to the same cell if and only if they agree on the meaning of 'Salad is delicious'. The *meaning** relation determines another such partition, in which worlds belong to the same cell if and only if they agree on the *meaning** of 'Salad is delicious'. By positing that the meanings of semantic vocabulary shift more rapidly than those of non-semantic vocabulary, and shift in such a way as to make 'All worlds and times close to actuality and the present match actuality and the present as regards the meanings of non-semantic vocabulary' semantically express a truth at every time and world, one can block the conclusion that anyone ever goes wrong in homophonically reporting neighbouring utterances of sentences like 'Salad is delicious'.

The suggestion that semantic vocabulary is unusually shifty compared to non-semantic vocabulary might initially seem strange. But it can be made to seem more familiar by thinking of it as a view on which semantic vocabulary is importantly similar to words like 'actually' and 'now'. For of course the semantic contributions of these words do shift even over tiny modal and temporal distances: someone who said 'It is raining now' a few seconds ago did not thereby assert that it is raining *now*. Grouping

speech reports with 'actually' and 'now' will seem natural if one analyses them as involving some kind of covert indexical element. And such analyses are not unprecedented: consider, for example, Hartry Field's proposal that 'to say that a word means <rabbit> is just to say that it has the same meaning-characteristics as the meaning characteristics of my actual use of "rabbit"' (Field 2001: 159).

So, the idea that semantic vocabulary is shiftier than other vocabulary has some attractions, and by invoking it, patchwork theorists can avoid having to posit any errors in homophonic reports like Fred's, which involve only non-semantic vocabulary. But the problematic errors will come back in full force when we turn to reports of speeches involving semantic vocabulary. Suppose that at t_3 , shortly before the last word in

(13) Mary just said that someone deserved a vacation

acquired its current meaning, it was uttered by Susan. (If semantic vocabulary is unusually shifty, then plausibly all the words in (13) except for 'said' already had their current meanings at t_3 : in that case, what Susan asserted in uttering (13) was that Mary had just said* that someone deserved a vacation, where *saying** is whatever 'said' semantically expressed at t_3 .) At t_4 , shortly *after* the last word in (13) had acquired its current meaning, Frank uttered

(14) Someone just said that Mary just said that someone deserved a vacation

Given the compositional picture we have been taking for granted, in which shifts in the meanings of sentences are due to shifts in the meanings of words, we cannot save this speech of Frank's in the way we contemplated saving Fred's earlier speech. For every word in (14) also occurs in (13), and by stipulation, all the words in (13), including 'said', already had their current meanings at t_4 . Thus there is no escaping the conclusion that in uttering (14), Frank said that someone had just said that Mary had just said that

someone deserved a vacation. And if so, we are forced to admit that Frank said something false, since the only relevant utterance is Susan's utterance of (13), and in uttering (13), Susan did *not* say that Mary had just said that someone deserved a vacation.⁴²

Having spotted the flaw in our original argument about Fred and Sally, one might have hoped to vindicate some general principle to the effect that, setting aside a small class of indexical expressions, homophonic reports are true whenever the reporter is very similar to the subject. But this is a mirage. At best, by positing that semantic vocabulary undergoes semantic shifts more easily and frequently than other vocabulary, we can endorse a restricted version of the principle that applies only to reports of non-semantic speeches.⁴³

⁴² As before, some of the problematic reports could be saved by postulating that the relevant shifty vocabulary is subject to the phenomenon of 'inheritance' discussed in §3.3. In the present case, the suggestion would be that because of Frank's backward-looking focus, one or both of his uses of 'say' inherit the content of Susan's, so that his report is true. But for the reasons discussed in §3.3, we doubt that such techniques are general enough to do all the required work.

⁴³ We have been operating within a standard framework in which speech reports are thought of as describing relations between speakers and propositions. Some philosophers reject the ideology of propositions in favour of analyses on which the truth of speech reports turns on relations between speakers and sentence-tokens, such as Davidson's famous proposal that an utterance of 'A said that P' is true just in case the referent of 'A' produced a sentence-token that bears a certain relation called 'samesaying' to the component token of 'P' (Davidson 1968). One might think that this model would bring new resources to bear on our puzzle, especially given the option of denying the transitivity of samesaying. But in fact, much of the dialectic of the present section will replay itself in this new theoretical setting. Suppose that A, B and C respectively produce tokens of 'Salad is delicious', 'A said that salad is delicious' and 'B said that A said that salad is delicious and A didn't say that salad is delicious', in such a way that B's token of 'Salad is delicious' samesays A's, and C's two tokens samesay B's token and one another, but do not samesay A's token. If we suppose that C's tokens of 'A said that salad is delicious' samesay B's, we will have to categorise C's utterance as true, and thereby admit that B has fallen into an error analogous to Fred's error in the main text. (This is an internal tension in the theory, since on the assumption that it applies to both C's and B's uses of 'say', it will classify B's utterance as true alongside C's error-ascribing utterance.) We can escape from this result by denying that the samesaying relation holds between B's and C's tokens of 'A said that salad is delicious', so that the first conjunct of C's speech is false: this is analogous to

4.2 *Epistemic worries*

Besides being uncomfortable in its own right, the conclusion that we sometimes go wrong in making clear-sighted homophonic reports of close neighbours also prompts an epistemic worry. When the underlying changes are gradual, there are no tell-tale signs which could enable us to distinguish cases of recent or imminent semantic change from cases of semantic stasis. How, then, can we *know* the propositions expressed by our cross-temporal reports?

While it is easy to make this epistemic problem sound gripping, some will think that it can be adequately addressed by relying on familiar general anti-sceptical moves. Why should the fact that the method of homophonic interpretation occasionally leads people astray preclude it from elsewhere delivering knowledge, any more than any other fallible method? One could bolster the judgment that we do have knowledge in ordinary cases by appeal to familiar safety- or reliability-based pictures of knowledge. On the other hand, such pictures can be taken too far in brushing aside epistemic dangers. Plausibly, it is at most a matter of decades until the next shift in the meaning of 'salad' (or in the meaning of 'say') leads to an outbreak of false homophonic judgments about recent speeches. Is this enough of an epistemic buffer? If aliens are in fact going to invade the Earth in 2050, it would be rather tendentious to claim that we can still know that they will not invade in the next year.

Even sharper problems arise if one adopts the natural assumption that different words shift their meanings at different times. On this view, it is probable that some words have undergone very recent semantic shifts, so that the method of reporting

the shifty view about semantic vocabulary considered in the main text. Indeed, given that samesaying theorists think of B's and C's utterances as involving demonstratives with different denotations, this is a rather natural move: the samesaying theory is another example of a theory that finds a concealed indexical element in speech reporting. However, the cost of this move is that homophonic reports of reports will sometimes be false even when the reporter and the subject are temporally or modally very close.

recent speeches homophonically is even now generating errors of the relevant sort. Thus, it is far from clear that the appeal to considerations of safety and reliability is adequate to counter the sceptical worries raised by the patchwork model.

A second response to the epistemic worries is to downplay the importance of knowledge, and look instead to vindicate rational confidence in the relevant reports. Suppose there is good reason to think that at least fifty years go by between shifts in the meaning of 'salad'; then even without further relevant evidence, we could reasonably have a credence of at least 98% that the practice of using 'salad' to mean *salad* was already up and running last year. The problem with resting content with high confidence is that the practice we are trying to justify involves *assertions* of the relevant reports. And as lottery cases show, high confidence is not always enough to make an assertion acceptable. *Something* goes wrong when people assert that they will lose the lottery just on the grounds that theirs is only one among many tickets. We think that bringing in the concept of knowledge to explain this datum is fruitful, but in any case, simply vindicating high confidence is insufficient to address the depth of the epistemic worry. Moreover, given the plausible view that different words shift at different times, it is hard to see how one could justify high confidence when it comes to a homophonic report of even a medium-length speech made last year.

In sum, we are not convinced that routine anti-sceptical manoeuvres suffice to dispel the epistemic worry for the patchwork model. However, developing these worries into a fully-fledged argument would require a more developed account of knowledge and assertion than the present paper can provide.

4.3 *Vagueness and the argument for Parity*

Let us now revisit the central motivation for *Parity* as developed in §2, namely that because the members of the relevant sets of candidate contents are so similar in relevant

underlying respects, it would be unacceptably arbitrary to hold that there are big differences as regards how easily they can get to be meanings or objects of speech acts.

When a philosophical view is faced with an objection based on arbitrariness, often the right response is to invoke vagueness. For example, if you are tempted to identify numbers with sets and run up against Benacerraf's (1965) objection that any particular such identification is arbitrary, you can try a response that says that while it is definitely true that all numbers are sets, there is no particular set that is definitely a number. Different views about the nature of vagueness will provide different interpretations of the 'definitely' operator that is used in expressing this claim. But provided that the background view of vagueness does not license the inference from 'Definitely something is F' to 'Something is definitely F', the anti-Benacerraf move will at least be a formally coherent option.

Let us see how this strategy looks when applied to a typical instance of *Parity*. The *Parity*-denier says that some members of the family Θ of propositions about the distance between Cup A and Cup B are unlike others in being asserted at positive-measure sets of worlds. In reply to the accusation that this makes an arbitrary division, the *Parity*-denier can say that 'assert' is vague in such a way that no proposition within Θ is either *definitely* asserted at a positive-measure set of worlds, or *definitely not* asserted at a positive-measure set of worlds. In general, the idea will be that when there are no physical or metaphysical discontinuities within a family of candidate contents, then they will all be alike as regards which semantic privileges they *definitely* enjoy and fail to enjoy.⁴⁴

⁴⁴ This suggestion is consistent with an epistemicist interpretation of 'Definitely P' as 'There is no obstacle of such-and-such sort to our knowing that P'. However, it is not clear how *satisfying* the response to the arbitrariness-based motivation for *Parity* is on this construal. Insofar as there is any force to the thought that it would be objectionable to posit big semantic disparities between contents that are metaphysically and physically

Is this strategy really a coherent option for deniers of *Parity*? Remember that the primary motivation for denying *Parity* is given by the argument that *Plasticity* leads to pervasive error in ordinary counterfactual speech reports like (2). Those who deny *Parity* on these grounds will thus be willing to assert that (2) is acceptable, and that if the set of worlds where the proposition that salad is delicious is asserted has zero measure, (2) is not acceptable. But in advancing these premises, they are presumably committed to their definite truth: *definitely* (2) is acceptable, and *definitely*, if the set of worlds where the proposition that salad is delicious is asserted has zero measure, (2) is not acceptable. It follows from these claims that definitely, it is not the case that the set of worlds where the proposition that salad is delicious is asserted has zero measure. But the central idea of the strategy we are considering involves claiming that, within certain continuum-sized sets of propositions that are on a par in relevant underlying respects, *no* proposition is definitely asserted at a positive-measure set of worlds. And the proposition that salad is delicious belongs to such a set if any proposition does. Isn't this a contradictory collection of commitments?

Perhaps not. The motivation for denying *Parity* requires us to accept (15):

- (15) Definitely, the proposition that salad is delicious is asserted at a positive-measure set of worlds.

And the vagueness-based way of resisting the arbitrariness argument for *Parity* requires us to accept (16):

- (16) The proposition that salad is delicious is not such that it is definitely asserted at a positive-measure set of worlds.

But it is not obvious that (15) and (16) are contradictory, since it is not obvious that the inference-form 'Definitely, the proposition that P is F; therefore, the proposition that P

similar, its force does not seem to be diminished by the further claim there is an obstacle of a certain sort to our *knowing* which contents enjoy the relevant semantic privileges.

is such that definitely, it is F' is valid. This inference form can be thought of as expressing the "rigidity" of terms of the form 'the proposition that P' with respect to the 'definitely' operator. But while such terms are uncontroversially rigid in the usual modal sense, there are many theories of vagueness according to which modal rigidity is not sufficient for rigidity with respect to 'definitely'. For example, many hold that even though it is definite that Kilimanjaro is a mountain, nothing is definitely a mountain (see McGee and McLaughlin 2000). Such non-rigid behaviour in ordinary names might be extended to terms of the form 'the proposition that P'. For example, one might claim that definitely, the proposition that Kilimanjaro is tall is a singular proposition whose subject is a mountain, even though there is no proposition that is definitely a singular proposition whose subject is a mountain. Those who hold that every proposition is either definitely true or definitely false (e.g. Dorr 2003) will certainly need to resist rigidity in this way. By contrast, it would be natural, though perhaps not inevitable, for those who believe in vague propositions as well as vague sentences to accept the rigidity inference.⁴⁵ The fact that the present strategy for resisting the arbitrariness-based argument for *Parity* requires an account of vagueness that does not license the rigidity inference thus significantly restricts its dialectical effectiveness.⁴⁶

4.4 *Vagueness about the location of patch boundaries*

So far, the ideology of vagueness looks like a helpful resource for proponents of the patchwork model. However, the situation becomes more complicated when we turn to

⁴⁵ The most common motivation for accepting vague propositions is a general commitment to the equivalence of 'Definitely P' and 'The proposition that P is definitely true'. Those who accept this equivalence must accept instances of the rigidity inference in which 'F' is replaced by 'true' or 'not true', assuming that they also accept the schema 'Definitely (P if and only if the proposition that P is true)'. It is unclear whether a principled account could be developed on which the rigidity inference sometimes fails for other values of 'F', but the factors which lead to these failures never make for failures involving 'true' or 'not true'.

⁴⁶ Similar issues arise for 'the property of being ϕ ', 'the relation of ϕ -ing', etc.

the worry that drives the arbitrary-boundaries argument—the worry that no principled answer can be given to the question why the modal and temporal boundaries between semantically homogeneous patches should be where they are. Taken by itself, this is too Sorites-like to constitute a compelling objection to the patchwork model. But once we bring the ideology of vagueness into play, a more formidable version of the objection comes into view.

Suppose the microphysical facts relevant to the meaning of 'salad' have been changing at a roughly constant rate, so that there are no principled differences between one-hour intervals which we could point to explain why 'salad' should have shifted its meaning during this one-hour interval rather than that one. Even if the similarity of the one-hour intervals fails to support the conclusion that if any of them contain semantic shifts all of them do, it might support the weaker conclusion that if any of them contain semantic shifts, none of them *definitely* fails to contain a semantic shift. Vagueness is ubiquitous, in semantic vocabulary as everywhere else. If certain items (in this case, hour-long intervals) are very similar in all the respects to which our use of a certain predicate (in this case, 'contains a shift in the meaning of "salad"') is sensitive, the default conclusion is that either all the items are definitely in the predicate's extension, or all the items are definitely not in the predicate's extension, or all the items are borderline cases of the predicate.

But given the patchwork model, the conclusion that no one-hour interval definitely fails to contain a shift in the meaning of 'salad' seems disastrous. Suppose it is an hour since Sally uttered 'Salad is delicious'. If it is indefinite whether 'salad' has shifted its meaning since then, then it is indefinite whether Sally said that salad is delicious. If you utter 'Sally said that salad is delicious', your speech will be bad in whatever way it is bad to assert 'Harry is bald' when Harry is only a borderline case of baldness. Whatever

the right theoretical account of this form of badness might be, the conclusion that it is pervasive in our cross-temporal and counterfactual speech reports seems little better than the conclusion that such reports are pervasively false.⁴⁷

While this new and improved version of the arbitrary boundaries argument—call it the *vague boundaries argument*—is not just a rehash of the Sorites, the issues it raises do arise for many non-semantic subject matters. In particular, they arise in many cases when predicates with the logic of equivalence relations apply over continuously varying domains. Consider, for example, ‘equally funny’, as applied to a series of possible situations stretching from a paradigmatically funny situation to a paradigmatically unfunny one, in which each situation differs from its neighbours to the same minute degree in each relevant underlying respect. While it is tempting to apply ‘equally funny’ to any particular pair of neighbouring situations, we cannot just flatly assert that *all* neighbouring pairs are equally funny, since ‘equally funny’ has the logic of an equivalence relation. Confronted with the entire sequence, there is also a temptation to say that *none* of the neighbouring pairs are equally funny; but this move is also very radical, since it suggests that nearly all ordinary applications of ‘equally funny’ are in error. Thus we seem forced to think of the series as divided into cells of equal funniness, so that some neighbouring pairs straddle cell-boundaries while most do not. But when we bring the ideology of definiteness into play, a new challenge rears its head. Given that all the pairs of neighbours are on a par in the respects to which our use of ‘equally funny’ is sensitive, we seem forced to say that no two neighbours are definitely equally funny, or definitely not equally funny. And this entails that almost all ordinary

⁴⁷ Note that the sense of ‘definitely’ in play is a specifically vagueness-theoretic one, on which ‘*a* is definitely *F*’ is equivalent to ‘*a* is *F* and *a* is not a borderline case of *F*-ness’. So it does not help for the proponent of the patchwork view to point to some *other* interpretation of ‘definitely’ under which it can be reasonable to assert that *P* even when one knows that it is not definitely the case that *P* (e.g. an interpretation where ‘Definitely *P*’ means ‘The objective chance that *P* is 1’).

applications of 'equally funny' are problematic in the same way as applications of 'bald' to borderline-bald Harry.

A full survey of the possible responses to this argument would require a deeper investigation of the foundational issues raised by vagueness than we have room for here. Our view is that the best response will appeal to the thought that vagueness typically introduces a kind of context-sensitivity.⁴⁸ On particular occasions when a vague expression is used, its vagueness will often be partially "resolved". For example, by pointing to a certain borderline-red object and saying 'This is red', one can create a local setting in which one can legitimately apply 'red' to this and various other things, even though there are other settings in which it would have been equally legitimate to apply 'not red' to those things. Moreover, facts about the salience of certain similarities play an important role in determining how these local standards work: for example, when two things are saliently similar in colour, one can legitimately say 'If one of these things is red, both are'. One could likewise plausibly hold that whenever two situations are saliently similar in funniness-relevant respects, it will be legitimate to assert 'These situations are equally funny', even though it is also legitimate (in any context) to assert 'It is not the case that any two situations that are similar enough in funniness-relevant respects are equally funny'.

Can the same idea be used to rescue the patchwork model from the vague boundaries argument? As applied to that subject matter, the idea would be when one is reporting on a particular temporally or counterfactually close situation, the vagueness of semantic vocabulary is partly resolved so as to make the report legitimate. For example, when we are talking about Sally's utterance of 'Salad is delicious' an hour ago, precisifications of 'contains a shift in the meaning of "salad"' according to which the

⁴⁸ Among those who connect vagueness and context-sensitivity in this way are Lewis (1979), Raffman (1994) and Soames (1999). Fara (2000) defends related ideas.

most recent hour-long interval is in its extension are excluded from the range of precisifications that need to be taken into account in evaluating our homophonic report of Sally's speech.

By invoking this kind of contextualism about semantic vocabulary, the patchwork model can save a wide range of homophonic reports of speeches which do not themselves involve semantic vocabulary. But problems come back when we turn our attention to homophonic *reports of reports*, just as such reports made trouble in §4.1 for the view that semantic vocabulary is unusually plastic. The shifting of standards to which vagueness gives rise is a kind of context-sensitivity: utterances in which vagueness is resolved in different ways involve the assertion of different propositions, or at least do not definitely involve the assertion of the same proposition. Suppose that Fred uttered 'In August, Sally said that salad was delicious' once in September and once in October. On the second occasion, the similarity between the usage of 'salad' in August and October was salient to him. Because of this, the vagueness of semantic vocabulary was resolved on that occasion in such a way as to make 'Sally in August meant the same thing by "salad" that I mean now' legitimately assertable. During the September speech, the similarity between September and August was salient to Fred, but the similarity between *October* and August was not salient to him, since he was not thinking about October. On the proposed account, this difference in the salience facts means that the vagueness of 'say' is resolved differently in Fred's September and October speeches. Thus, it is not definitely the case that there is a single proposition that Fred asserted both in September and in October. So if you, on some later date, say 'Fred has twice said that in August, Sally said that salad was delicious', your speech will not be definitely true, even if the similarity between all the earlier dates is salient to you on that later occasion. No minor tweak will help: it is of the essence of the approach we are

considering that it buys us the right to discount certain small differences in use as irrelevant to the meaning of non-semantic vocabulary only at the price of making small differences extremely relevant in the case of semantic vocabulary.

In summary, while a contextualist treatment of the vagueness of cell-boundaries provides a satisfying story about predicates like 'equally funny', its capacity to protect the patchwork model from the vague boundaries argument is considerably more limited.

4.5 *Vague boundaries and degrees of truth*

There is a different, non-contextualist strategy for getting around the problem of vague cell-boundaries for 'equally funny' which, if it worked, might allow proponents of the patchwork model to resist the vague boundaries argument without reintroducing new problems about reports of reports. This strategy appeals to the idea that vague sentences enjoy varying 'degrees of truth' between 0 and 1, such that the acceptability of a speech does not require the sentence uttered to have a degree of truth quite as high as 1. While the most familiar view using this ideology involves a thorough rejection of classical logic and hence goes outside our stated purview, there are versions that are more conservative in this respect. For example, by imposing a measure on the set of admissible precisifications of the language, one can characterize a sentence's degree of truth as the proportion of admissible precisifications on which it is true (see Lewis 1970, p. 229 and Kamp 1975). Given any pair of neighbouring situations, the set of precisifications which take that pair to mark a cell boundary for 'equally funny' will have a very small measure compared to the set of precisifications which don't; thus it is easy to secure degrees of truth close to 1 for all attributions of 'equally funny' to neighbouring pairs. An analogous account can be given for semantic vocabulary: when the modal and/or temporal distance between a given homophonic report and its target is small, most of the admissible ways of drawing the cell-boundaries will place the

reporter and the subject in the same cell, and thus the report can be assigned a degree of truth very close to 1.

One hallmark of this classical-logic-friendly degree theory is that there are situations where several sentences individually have high degrees of truth, while the negation of their conjunction also has a high degree of truth. The general strategy for proponents of the account will be to assimilate this phenomenon to the preface paradox—another case where it may seem not so bad to assert a number of propositions while also asserting the negation of their conjunction. In fact, when we set aside vagueness-theoretic jargon, the pattern of speech that we find from proponents of the patchwork model who regard the facts about the locations of the patch-boundaries as vague and take high (expected) degree of truth as their guide to assertion will be close to the pattern of speech of those who are uncertain about the location of the cell-boundaries, and are willing to assert that *P* whenever they have a high credence that *P*.

A subtler, and we think more plausible, way to connect degrees of truth with assertoric utterance would require a high degree of truth for 'I know that Φ ' as a condition for Φ to be felicitously asserted. There is no obvious reason why Φ and 'I know that Φ ' couldn't both have degrees of truth close to but not equal to 1. (For example, if the cells are big according to each precisification, it might be that on most of them, the reporter and the subject are *safely* far from the cell-boundaries.) Those who regulate their assertions in this way will also be willing to assert many sentences individually while asserting the negation of their conjunction. In their speech patterns, they will sound like people who have a high credence that they *know* that any given pair of close situations is not separated by a cell-boundary, and are willing to assert that *P* whenever they have a high credence that they know that *P*.

We have three critical remarks to make about the use of degrees of truth to save the patchwork model. First, there is a worry about longer reports. Even if most precisifications of semantic vocabulary agree that any given word has not changed its meaning since last year, for an indirect report of a medium-length speech made last year to have a high degree of truth, most precisifications would have to agree that *none* of the relevant words have changed their meaning since last year. This could be achieved if, according to most precisifications, different words tend to change their meanings at the same time, so that most one-year periods involve no semantic changes at all. But there is something unattractive about positing such synchronization, even when it is granted that the question when these sweeping semantic changes occur has no definitely true answer.

Second, we have expressed some sympathy for the idea that we cannot know, e.g., that people last year thought that salad was good for them, if the method on which we base our belief is one that has fairly recently generated false beliefs. If this is right, then since most admissible precisifications will have someone making a mistake in the not-so-distant past or future, the degree of truth of a claim like 'I know that people last year thought that salad was delicious' will be low. On the more plausible knowledge version of the view, the low degree of truth of such knowledge reports will undermine the assertability of ordinary cross-temporal reports.

Third, insofar as one's epistemic filter for assertoric utterance is high (expected) degree of truth (of the sentence uttered or the corresponding knowledge claim), one's pattern of discourse will risk being blind to the logical relationships of one's current speech to earlier segments of the discourse. Typically, having asserted some sentences individually, one has to somehow signal that one is conversationally switching gears before asserting something logically inconsistent with their conjunction. Views on

which vagueness involves context-sensitivity have a natural strategy for explaining this phenomenon: facts about what has already been said in a conversation will be among the factors that affect how vagueness is resolved. By contrast, no natural strategy for explaining the relevance of prior discourse is forthcoming from the degrees-of-truth view.

A complete appraisal of the ideology of degrees of truth would require a more extensive discussion of its foundations than is appropriate here. Our own feeling is that an adequate foundational account of the puzzles of vagueness will not have much use for the ideology, and will confirm our view that the vagueness-based argument of §4.4 presents a serious additional problem for the patchwork model of semantic shifts.

5. Rejecting *Finitude*: the multiple propositions view

Finally, we turn to the strategy for resisting the continuum argument that endorses *Parity* while denying *Finitude* and *Plasticity*. Consider the example from §2, where W is the set of possible worlds with the same total energy as the actual world, and Γ is the set of propositions to the effect that we assert p , for each proposition p within a certain continuum-sized set Θ of close variants of the proposition that Cup A is pretty far Cup B. To deny *Finitude* in this case is to claim that infinitely many members of Θ are asserted at some worlds in W .⁴⁹ To take this suggestion seriously, one would need to think that it was fairly easy for people like us to assert infinitely many propositions by making a single assertive utterance. Indeed, the most promising version of the view would seem to be one on which *almost all* assertoric utterances involve the assertion of

⁴⁹ As noted in footnote 15, blocking the argument from *Parity* to *Plasticity* actually requires the stronger claim that the set of worlds at which infinitely many members of Θ are true has positive measure.

infinitely many propositions, and similarly for other speech-acts.⁵⁰ Call this claim 'Multiplicity'.⁵¹

When Γ is a set of propositions, let Γ^* be the set of all Γ -profiles, where a Γ -profile is any conjunction of some propositions in Γ with the negations of all the others. We can use Γ^* as the basis for a new instance of the continuum argument. Insofar as considerations of arbitrariness supported *Parity* for Γ , they will almost certainly support *Parity* for Γ^* . For example, the hypothesis that countably many *subsets* of Θ are distinguished in being such that the proposition that we assert each of them and no other member of Θ has positive measure seems just as repugnant as the hypothesis that countably many *members* of Θ are distinguished in being such that the proposition that we assert them has positive measure. But since any two members of Γ^* are logically inconsistent, denying *Finitude* for Γ^* is out of the question. Those who resisted *Plasticity* for Γ by denying *Finitude* are thus under pressure to embrace *Plasticity* for Γ^* . We will therefore focus on a version of Multiplicity according to which small differences between worlds typically induce small differences in the complete profile of propositions that are the objects of speech-acts, so that for any collection of propositions of the sort we are concerned with, the set of worlds where all and only the members of that collection are asserted has zero measure. In principle one could endorse the plasticity of "assertion profiles" across worlds but not across times; but for ease of

⁵⁰ The idea that we often say multiple things in making a single utterance has recently been defended by Soames (2002), Cappelen and Lepore (2005), and Cappelen (2008). These authors do not, however, commit themselves to the number ever being infinite, or to the claim that when one asserts a proposition one typically asserts many *very similar* propositions. In the realm of mental content, Lormand (1996) argues for a 'multiple meanings view' that has much in common with Multiplicity.

⁵¹ The idea that multiple propositions are asserted by a single utterance could also be developed in such a way as to respect *Finitude*. While this package does not have the advantage of saving *Parity*, one might still hope that it would be useful in addressing the problem of occasional error in homophonic reporting discussed in §4.1. Much of the discussion that follows concerning issues of error for certain versions of Multiplicity will carry over to these milder multiple-assertion views.

exposition, we will assume that the proponent of Multiplicity takes the same view of the temporal case. That is, even within a world, it almost never happens that two utterances of a sentence like 'Salad is delicious' involve the assertion of *exactly* the same range of propositions, although there is often substantial overlap (when the temporal distance is not too great).

As we saw in §3 and §4, both the *Plasticity*-embracing and *Parity*-denying approaches face significant trouble with counterfactual and cross-temporal homophonic reports. How does Multiplicity fare in this respect?

There are some ways of developing Multiplicity that are useless for the purposes of vindicating such reports. Consider, for example, the view that in uttering a sentence assertorically, one asserts all and only the propositions *entailed* by the unique proposition that the sentence semantically expresses relative to the context. Implausible though it is, some theorists do talk as if assertion were closed under entailment in this way. But plausible or not, this approach is only marginally useful in saving our reporting practice. It allows a homophonic report to be true when the semantic content of the embedded sentence at the target time and world is *stronger* than its semantic content of that sentence at the time and world of the report. But in any other case, it will entail that the reporter asserts the false proposition that the subject asserted *p*, where *p* is the proposition semantically expressed by the embedded sentence relative to the reporter's context. (For a term like 'salad', since there are so many dimensions of potential variation, it would be unsurprising if most shifts involve strengthening along some dimensions and weakening along others.) We take it that this result is problematic: if one is sincere, one doesn't assert things one doesn't believe, and nothing about Multiplicity makes the conclusion that ordinary people are constantly forming false beliefs about one another's speech acts any more palatable.

Analogous problems arise for many other ways of developing Multiplicity. Consider, for example, a view according to which an assertoric utterance counts as an assertion of every proposition that is 'close enough' to the (unique) proposition semantically expressed by the sentence uttered. Suppose one develops the relevant notion of closeness in such a way that whenever p is close enough to q , the proposition that x asserted p is close enough to the proposition that x asserted q . If semantic expression behaves in a modally/temporally plastic way, then the range of close propositions will vary correspondingly: with each small change in the semantic content, there will be a small shift in the range of propositions close to the semantic content. As a result of this, it will be typical for homophonic reports to involve the assertion of falsehoods: whenever p is close to the proposition now expressed by Φ but not close to the proposition expressed by Φ at an earlier time t , a present utterance of 'Someone asserted that Φ at t ' will involve the assertion of the false proposition that someone asserted p at t .⁵²

As a third example, consider a version of Multiplicity partly inspired by standard supervaluationist semantic ideology.⁵³ A language has many "admissible interpretations", each of which maps sentences to propositions in accordance with standard compositional principles. Further—and this is where the view goes beyond standard supervaluationism—when one literally utters a declarative sentence, one

⁵² In §3.4 we discussed the 'close enough proposition' idea as an account of *quasi-assertion*, the relation that figures in the "loose content" of counterfactual reports using 'assert'. The objection just presented can be adapted to this view. Suppose that, if Fred had been here, he would have uttered 'Sally asserted that salad is delicious' and thereby quasi-asserted a false proposition. Then the loose content of 'If Fred had been here, he wouldn't have asserted anything false' is false, even though its strict content may be true. Given that honest people believe the loose contents of the sentences they utter assertively, ordinary assertive utterances of sentences like 'If Fred had been here, he wouldn't have asserted anything false' will betray the kinds of pervasive false beliefs that the 'loose talk' idea was designed to avoid.

⁵³ As Williamson (1994) points out, this ideology is in itself perfectly neutral as regards the nature of vagueness.

thereby asserts all and only the propositions to which that sentence is mapped by admissible interpretations. This view also predicts that homophonic reporting typically involves the assertion of falsehoods. For consider a proposition p that is among the propositions to which 'Salad is delicious' is mapped by admissible interpretations at t_2 , but not at t_1 . Suppose Sally utters 'Salad is delicious' at t_1 and Fred utters 'Sally asserted that salad is delicious' at t_2 . And suppose that there is at least one admissible interpretation of Fred's language which maps 'Sally' to Sally, 'assert' to the relation of assertion and 'Salad is delicious' to p . Then one of the propositions that Fred asserts is the proposition that Sally asserted p . And this proposition is false, since at t_1 no admissible interpretation mapped 'Salad is delicious' to p . Of course, Fred may also have asserted many true propositions. But we don't see how this does anything to mitigate the implausibility of the claim that he asserted something false.⁵⁴

The problems with the three versions of Multiplicity we have just considered come from the mechanical way in which they attempt to determine the range of propositions

⁵⁴ Could we save Fred's speech by positing a "penumbral connection" between 'salad' and 'assert', whereby the admissible interpretations at t_2 that map 'Salad is delicious' to p also map 'assert' to relations other than assertion, which Sally did bear to p ? We doubt that there is a recipe for a general solution here. One way to bring this out is to consider a case where semantic expressions like 'assert' and 'admissible interpretation' are the only relevant semantically shift items. Consider the following argument:

These days, people who use 'assert' thereby make assertions about assertion. But this has not always been the case, since the use of 'assert' has been evolving gradually from some very different starting point. There must thus have been a first minute during which people using 'assert' began making assertions about assertion. Now, consider an utterance at that time of 'People who used "assert" five minutes ago made assertions about assertion'. One of the propositions asserted in that utterance was the false proposition that people five minutes before then made assertions about assertion. So some people in our history are making a mistake when they use homophonic reporting methods even across tiny temporal intervals.

Assuming that 'Admissible interpretations are compositional, and one asserts all the propositions to which one's sentence is mapped by admissible interpretations' is itself true on all admissible interpretations, this argument will be sound on all admissible interpretations. Thus the proposed penumbral connections cannot save homophonic reports involving 'assert'. And as we said in §4.1, we find it hard to take seriously views that endorse such reporting for 'salad' but engages in special pleading for 'assert'.

asserted by utterances of a given sentence. As soon as a proposition starts to be asserted by unembedded utterances of a sentence, propositions about that proposition start to be asserted by reports which embed that sentence under attitude and speech-act verbs.⁵⁵ Because of this, various cross-temporal and counterfactual reports end up involving the assertion of false propositions. For Multiplicity to help, it needs to be implemented in a more flexible way.

Here is one toy view that would certainly help. Suppose that there is a hierarchy among the set of admissible interpretations of the language: some are “level 1”, some are “level 2”, etc. Over time, an interpretation that begins to be admissible will start at admissibility level 1, rise up for a while, then sink back down to 1 before finally ceasing to be admissible. When one utters a sentence sincerely, one thereby counts as asserting all and only the propositions to which it is mapped by admissible interpretations whose level exceeds the sentence’s *embedding depth*. ‘Embedding depth’ is defined as follows: sentences with no attitude or speech act verbs have embedding depth 0; sentences in which all the complements of such verbs have embedding depth 0 have embedding depth 1; and so on. Returning to our example, when Fred at t_2 says ‘Sally asserted that salad is delicious’, since that sentence has embedding depth 1, he asserts all and only those propositions to which it is mapped by interpretations whose admissibility level at t_2 is at least 2. Presumably, at t_1 , ‘Sally is asserting that salad is delicious’ was true on *all* admissible interpretations (i.e. interpretations whose admissibility level was at least 1). But since t_2 and t_1 are close, it is plausible that every interpretation whose level at t_2 is 2 or more was already admissible at t_1 . Hence each of the propositions Fred asserted at t_2 by uttering ‘Sally asserted that salad is delicious’ is true.

⁵⁵ The version of the supervaluationist approach discussed in note 48 avoids this feature for ‘salad’ but not for ‘assert’.

This implementation of Multiplicity, like those considered earlier, proposes a rule whereby the structure of a sentence determines the set of propositions asserted by uttering that sentence. While it does avoid the problem confronting the earlier versions of Multiplicity, the method by which it does so is highly artificial. Arguably, those who want to solve the puzzle by appeal to Multiplicity would do better to jettison the idea that there is a systematic recipe whereby the facts about what is asserted can be read off the semantic profile of the sentence uttered. We could keep the ideology of a range of admissible interpretations each of which obeys compositional principles, while thinking of the facts about these interpretations as merely *constraining* what speakers can use a sentence to assert, rather than determining it. There will often be substantial differences in plausibility within the set of propositions to which admissible interpretations map a given sentence; it is attractive to suppose that speakers who are appropriately sensitive to these differences (and who expect their audiences to be similarly sensitive) will be able to utter the sentence while only asserting some of the more plausible members of that set.

To flesh out this proposal, we will need a story about the nature of assertion and belief which allows us to make sense of the required fine-grained sensitivity. We will have to be careful to avoid conceptions of belief that re-enact the problems we have been developing for assertion. For example, we can have nothing much like the popular picture on which to believe a proposition is to for it to be semantically expressed by a mental sentence tokened in one's belief box, where the relevant relation of semantic expressing works compositionally. On such a view, anyone whose belief box contains a sentence like 'I have believed that Φ for a while' will believe some false proposition to the effect that he or she has believed p for a while, where p is some proposition that has only just now begun to be expressed by Φ . The challenge for friends of Multiplicity is

to tell a foundational story about attitudes and speech acts whose architecture is sufficiently different from the belief box picture that no parallel puzzle can arise.

6. Conclusion

In view of the elementary character of the considerations needed to get the puzzle of semantic plasticity going, it is rather surprising that it has not received attention among analytic philosophers. Perhaps it initially looked like a sophisticated trick, amenable to solution by appeal to context-sensitivity, vagueness, loose talk, or some other standard resource from the philosopher's toolkit. But as we have seen, there are no easy answers. We hope that the initial survey of options which we have undertaken in this paper will spur further interest and exploration.⁵⁶

⁵⁶ Thanks to Suzanne Bobzein, Jeremy Goodman, Anandi Hattiangadi, Ofra Magidor, Josh Parsons, David Sedley, and Timothy Williamson, to audiences at NYU, Oxford, Leeds and Nottingham, and to two anonymous referees who gave us very helpful comments.

Works cited

- Albert, David Z. (2000). *Time and Chance*. Cambridge, Mass.: Harvard University Press.
- Arntzenius, Frank (2006). 'Infinity, Relativity and Smoothness'. In John Hawthorne (ed.), *Philosophical Perspectives 20: Metaphysics* (Oxford: Blackwell): 1–16.
- and Cian Dorr (2012). 'Calculus as Geometry'. In Frank Arntzenius, *Space, Time and Stuff* (Oxford: OUP).
- Benacerraf, Paul (1965). 'What Numbers Could Not Be'. *Philosophical Review* **74**: 47–73.
- Bennett, Karen and Dean Zimmerman (eds.) (forthcoming). *Oxford Studies in Metaphysics* vol. 8. Oxford: OUP.
- Burge, Tyler (1979). 'Individualism and the Mental'. *Midwest Studies In Philosophy* **4**: 73–121.
- Cappelen, Herman (2008). 'Content Relativism and Semantic Blindness'. In Manuel Garcia-Carpintero and Max Kölbel (eds.), *Relative Truth* (Oxford: OUP): 265–86.
- and John Hawthorne (2009). *Relativism and Monadic Truth*. Oxford: OUP.
- and Ernest Lepore (2005). *Insensitive Semantics: A Defense of Semantic Minimalism and Speech Act Pluralism*. Malden, Mass.: Blackwell.
- Dasgupta, Shamik (forthcoming). 'Absolutism vs Comparativism About Quantities'. In Bennett and Zimmerman forthcoming.
- Davidson, Donald (1968). 'On Saying That'. *Synthese* **19**: 130–46.
- Dorr, Cian (2003). 'Vagueness Without Ignorance'. In John Hawthorne and Dean Zimmerman (eds.), *Philosophical Perspectives 17: Language and Philosophical Linguistics* (Oxford: Blackwell): 83–113.
- and John Hawthorne (forthcoming). 'Naturalness'. In Bennett and Zimmerman forthcoming.
- Fara, Delia Graff (2000). 'Shifting Sands: An Interest-Relative Theory of Vagueness'. *Philosophical Topics* **28**: 45–81. Originally published under the name 'Delia Graff'.
- Field, Hartry (1980). *Science Without Numbers*. Oxford: Blackwell.
- (2001). 'Attributions of Meaning and Content'. In *Truth and the Absence of Fact* (Oxford: Clarendon): 157–74.

- von Fintel, Kai (2001). 'Counterfactuals in a Dynamic Context'. In *Ken Hale: A Life in Language*, ed. Michael Kenstowicz (Cambridge, Mass.: MIT Press): 123–52.
- Hájek, Alan (MS). *Most Counterfactuals are False*.
- Hattiangadi, Anandi (MS). 'Semantic Zombies'.
- Hawthorne, John (2006a). 'De Re Determinism'. In Hawthorne 2006b: 239–44.
- (2006b). *Metaphysical Essays*. Oxford: Clarendon.
- Jackman, Henry (1999). 'We Live Forwards but Understand Backwards: Linguistic Practices and Future Behavior'. *Pacific Philosophical Quarterly* **80**: 157–77.
- Kamp, Hans (1975). 'Two Theories About Adjectives'. In Edward L. Keenen (ed.), *Formal Semantics of Natural Language: Papers from a Colloquium Sponsored by the King's College Research Centre, Cambridge*. Cambridge: Cambridge University Press: 123–55.
- Kearns, Stephen and Ofra Magidor (2008). 'Epistemicism about Vagueness and Metalinguistic Safety'. In John Hawthorne (ed.), *Philosophical Perspectives 22: Philosophy of Language* (Oxford: Blackwell): 277–304.
- (2012). 'Semantic Sovereignty'. *Philosophy and Phenomenological Research* **85**: 322–50.
- Lewis, David (1970). 'General Semantics'. *Synthese* **22**: 18–67. Reprinted with postscripts in *Philosophical Papers, vol. 1* (Oxford: OUP, 1983): 189–232.
- (1973). *Counterfactuals*. Oxford: Blackwell.
- (1976). 'Probabilities of Conditionals and Conditional Probabilities'. *Philosophical Review* **85**: 297–315.
- (1979). 'Counterfactual Dependence and Time's Arrow'. *Noûs* **13**: 455–76. Reprinted with postscripts in *Philosophical Papers, vol. 2* (Oxford: OUP, 1986): 32–66.
- (1983). 'New Work for a Theory of Universals'. *Australasian Journal of Philosophy* **61**: 343–77.
- Loewer, Barry (2007). 'Counterfactuals and the Second Law'. In Huw Price and Richard Corry (eds.), *Causation, Physics, and the Constitution of Reality: Russell's Republic Revisited* (Oxford: OUP): 293–326.
- Lormand, Eric (1996). 'How to Be a Meaning Holist'. *Journal of Philosophy* **93**: 51–73.
- McGee, Vann and Brian McLaughlin (2000). 'The Lessons of the Many'. *Philosophical Topics* **28**: 129–51.
- Moss, Sarah (2013). 'Subjunctive Credences and Semantic Humility'. *Philosophy and Phenomenological Research* **87**: 251–78.

Partee, Barbara (1989). 'Binding Implicit Variables in Quantified Contexts'. In Caroline Wiltshire, Randolph Graczyk, and Bradley Music (eds.), *Proceedings of the Chicago Linguistics Society* 25, University of Chicago Press, Chicago, 342–65.

Raffman, Diana (1994). 'Vagueness Without Paradox'. *Philosophical Review* **103**: 41–74.

Soames, Scott (1999). *Understanding Truth*. Oxford: OUP.

— (2002). *Beyond Rigidity: The Unfinished Semantic Agenda of Naming and Necessity*. Oxford: OUP.

Stalnaker, Robert C. (1968). 'A Theory of Conditionals'. In Nicholas Rescher (ed.), *Studies in Logical Theory*: American Philosophical Quarterly Monograph Series, No. 2 (Oxford: Blackwell): 98–112.

Williams, J. Robert G. (2012). 'Counterfactual Triviality: A Lewis-Impossibility Argument for Counterfactuals'. *Philosophy and Phenomenological Research* **85**: 648–70.

Williamson, Timothy (1994). *Vagueness*. London: Routledge.

— (1997). 'Replies to Commentators'. *Philosophy and Phenomenological Research* **57**: 945–53.