METASEMANTICS AND OBJECTIVITY

Ori SIMCHEN*

If Jones is found guilty of negligence, and the verdict is never overturned, does it follow that the judgment that Jones was negligent is objectively true? If it does, then it is presumably a fact that Jones was negligent. What kind of fact is it? On the other hand, if it does not follow that the judgment is objectively true, what might still falsify it? Is there a "way things are" legally speaking that goes beyond actual judicial decisions upheld by the courts? Such questions as these are most often raised regarding specific domains of judgment, such as the domain of legality or of morality or of science. In this paper I propose to examine the issue of objectivity more generally, with the hope of shedding some light on domain-specific concerns.

Writers on objectivity typically set things up in the following oppositional way: Here we are with our X-type judgments. In order for such judgments to be true or false, there has to be something over there by virtue of which they are so-call it the "X-facts".[1] Now the question arises as to the nature of these X-facts: Are X-facts *really* there to be discovered by us, or are they actually here in some sense, constituted by us? Are they judgment-independent or judgment-dependent? (I will not even pretend to do justice to the plethora of alternative ways of framing this type of query). Once we have settled on the metaphysics of X-facts, on the approach being considered, we can then turn to the question of epistemic access to them.

By characterizing the relation between the thinker and the X-facts in such a way, we already open up a divide that any account of the bearing

---

* University of British Columbia, USA.

[1] The very idea of facts as truth-makers for judgments has been the topic of much heated controversy in contemporary philosophy. It has been called into question by Frege, Gödel, and Davidson, among others. For the most recent round in the debate, see Stephen Neale, *Facing Facts*, Oxford, Oxford UP, 2001.

of the world on our judgments will be hard-pressed to close satisfactorily. Yes, X-facts are there to be discovered by us to the extent that we can get the facts wrong; but, No, they are not *entirely* independent of us to the extent that they attest to our conceptual involvement. And so the familiar back-and-forth refinements can continue without apparent end in sight. The history of philosophy has offered ample illustrations of the moral that if the metaphysical and the epistemic are pried apart at the very outset of accounts of the objectivity of our judgments, then putting them back together again will prove to be a formidable task. In this paper I hope to make some headway towards resisting the common temptation to pry them apart at the very outset. The strategy I will employ is best characterized by fastening on an imagistic contrast. As against setting up a yawning gap between the thinker and the facts and then turning to ask the metaphysical question about the nature of those facts, followed by the epistemic question about the thinker's access to them, I begin by focusing on those aspects of contact between the thinker and the world through which content emerges.

My strategy for broaching objectivity will be metasemantic and will follow in the footsteps of the so-called new theory of reference. We begin by considering the thinker in her worldly surroundings. We then ask the metasemantic question: How do the thinker's terms happen to gain the content that they do? Any plausible answer here will allude to what the worldly surroundings of the thinker actually are and to what the thinker's overall epistemic situation actually is. But the metasemantic explanatory strategy does not begin by addressing the nature of X-facts independently of the thinker's epistemic access to them and then proceed, as a separate project as it were, to address (or to set aside, as the case may be) the issue of epistemic access. Rather, the metasemantic strategy takes its point of departure from the basic idea that terms typically have content. This is sometimes referred to as their "intentionality", or "aboutness". And it is the possibility of such endowment that can be shown to require objectivity. Or so I will argue.

The argument to follow has a distinctly transcendental, and so Kantian, flavor. But a rather more direct route runs back from it to the work of the later Wittgenstein.[2] I will try to show that the very possibil-

---

[2] As in passages peppered throughout Ludwig Wittgenstein, *Philosophical Investigations*, 3. ed., Oxford, Blackwell, 1953, such as the following: "Let us imagine a table

ity that terms in a given domain should have content, the very possibility that they should contribute to the truth-conditions of claims in which they partake, depends on there being a distinction between what is relevantly the case and what only seems to be the case. More specifically, content-determination will be claimed to depend on the existence of an objective measure of correctness in extension-determination. For example, that a term such as "negligence" should have content demands that there be an objective measure of relevant similarity to paradigmatic instances of negligence, a measure of similarity that must be capable of transcending what merely seems to be relevantly similar to paradigmatic instances of negligence. If this is correct, then we still face a choice. We can either affirm the requisite measure of objectivity, or else we can deny the possibility that our terms have content after all. But such is an inevitable feature of transcendental arguments. Kant, for example, offers an argument purporting to show that if experience is at all possible, then the objects of experience must conform in certain elaborate ways to our cognition rather than the other way around. Assuming that this argument is successful, it is still open to us either to affirm Kant's consequent or deny his antecedent and conclude that experience is not possible after all. We may thus think of Kant's effort in this area as purporting to illustrate the heavy price we incur by denying his consequent, his Copernican revolution. Similarly in this case, if the argument to the effect that the possibility of content depends in certain elaborate ways on the relevant objectivity is successful, then it is still open to us to deny objectivity. But the price of such denial is the denial of the possibility that our terms have content. And that is a heavy price indeed.

Our terms have content. How do they gain it?This basic metasemantic question has a *prima facie* intelligibility. In what follows I will consider

(something like a dictionary) that exists only in our imagination. A dictionary can be used to justify the translation of a word X by a word Y. But are we also to call it a justification if such a table is to be looked up only in the imagination? –'Well, yes; then it is a subjective justification.' – But justification consists in appealing to something independent. – 'But surely I can appeal from one memory to another. For example, I don't know if I have remembered the time of departure of a train right and to check it I call to mind how a page of the time-table looked. Isn't it the same here?' – No; for this process has got to produce a memory which is actually *correct*. If the mental image of the time-table could not itself be *tested* for correctness, how could it confirm the correctness of the first memory? (As if someone were to buy several copies of the morning paper to assure himself that what it said was true.) (§265).

the extent to which it can be respected in its own terms. There is a cer-
tain philosophical tradition that considers such a question to be misdi-
rected, ill framed, or beg fundamental issues of philosophical methodol-
ogy. This is a tradition that takes the facts of meaningfulness to derive
from the interpretative situation. It is a tradition that takes the basic ques-
tion in the area to be something along the following lines: What makes it
the case that terms have the content that they do? Donald Davidson
clearly falls within this camp, as does, albeit in a different way, David
Lewis. Both take their most immediate inspiration from the work of W.
V. Quine. On Davidson's view, that expressions should have their con-
tent is a matter of being interpretable in this way according to a suitable
"interpretative" Tarskian truth-definition for the linguistic corpus to
which they belong, given the attitudes that speakers are likely to have
in their actual surroundings. On Lewis's view, that expressions should
have their content is a matter of the existence of an eligible mapping that
so assigns contents to them, where eligibility is a matter constrained both
by the attitudes of speakers, appropriately interpreted in turn, and by
what the plurality of possible worlds is really like. For both thinkers, what
makes it the case that terms have their content is at bottom a matter of
how they are interpreted. Call this way of thinking about meaningfulness
"metasemantic interpretivism".

If we situate speakers in their environment and raise the philosophical
question of how it is that their terms come to have their content without
privileging the interpretative situation, we part company with the above
tradition.[3] Call the alternative approach "metasemantic productivism".
Unlike the philosophical query after the determinants of the semantic
state of meaningfulness (How is it that terms *have* their content?), the
question raised by the rival approach (How is it that terms *come to have*
their content?) targets the determinants of a process– the process of gain-
ing semantic content. To this alternative approach belong first and fore-
most the efforts of Keith Donnellan, Saul Kripke, and Hilary Putnam.

---

[3] The qualification 'philosophical' before 'question' is important because neither
Davidson nor Lewis would deny the existence of interesting empirical questions in the
general area of content-determination. On the other hand, the qualification should not be
taken as a tacit endorsement of a contentious analytic-synthetic distinction – as if philo-
sophical questions can be sharply distinguished from empirical ones. These matters are
far subtler than first impressions reveal and I cannot deal with them in a satisfactory way
in the scope of this paper.

Various current attempts to naturalize the mind a la Jerry Fodor or Fred Dretske also belong here, where what is sought is a naturalist reduction of the intentional to the non-intentional. But reductive naturalism is not the only option in metasemantics. We can acknowledge that no such reduction of the intentional to the non-intentional is available while making genuine metasemantic explanatory progress within the general framework of metasemantic productivism. To see how this might be so, it is helpful to consider an analogy.

The Russell of Knowledge by Acquaintance and Knowledge by Description asks how it is possible for anyone other than Bismarck to grasp a proposition about Bismarck.[4] Most of us have never met the man, nor would we recognize him had we encountered him. Worse still, even if we had met him in the past or were perceiving him at the very moment of thinking or talking about him, while knowing full well that the man in front of us is Bismarck, we would not thereby be acquainted with the man himself but only with how he appears. In short, by Russellian lights we have no direct epistemic access to such items as Bismarck. Yet it is precisely such direct epistemic access to each constituent of a proposition that is required, according to Russell, to be in a position to grasp it. So Russell's answer to the aboutness question regarding Bismarck is roughly this: There is no possibility for anyone other than Bismarck of grasping a proposition that has Bismarck himself as a constituent. One can only grasp a proposition each element of which is an object of one's acquaintance, and only one's own sense data, one's self, and various universals, qualify as objects of one's acquaintance. However, we can grasp various propositions that describe Bismarck, propositions that Bismarck himself uniquely satisfies and that are composed of elements with which we *are* acquainted. It is only a proposition of this second type that anyone other than Bismarck can expresses with the words 'Bismarck was an astute diplomat'. Aboutness regarding Bismarck for anyone other than Bismarck is, on this view, a species of satisfaction.

Let us set aside the question of whether or not Russell's theory is correct. Consider someone who objected to it on the following grounds: Rather than offer a genuine answer to how aboutness regarding Bismarck is possible, the theory merely pushes back the question of

---

[4] Bertrand Russell, "Knowledge by Acquaintance and Knowledge by Description", *Proceedings of the Aristotelian Society*, 11, 1910, pp. 108-128.

aboutness to aboutness regarding one's own sense data, one's self, and various universals – in short, to aboutness regarding items with which one is acquainted. But the theory offers us nothing at all when it comes to how we can entertain propositions about objects of acquaintance; how, for example, one can entertain a proposition that has oneself as a constituent. Given this glaring lack, so the complaint concludes, no explanatory progress has been made by Russell's theory after all.

But the complaint is misguided. Explanatory progress would have been made by Russell's theory – if only it were otherwise plausible. To claim that Russell's explanation of aboutness regarding the likes of Bismarck is unsuccessful because it has not succeeded in eliminating any trace of aboutness from the explanans is to set the bar of explanation far too high. If Russell's theory were plausible, it would succeed in explaining how aboutness regarding all things reduces to aboutness regarding items of our acquaintance. That would have been a significant explanatory achievement. As it happens, the theory has little to recommend it on other grounds. But an account of aboutness need not culminate in a reduction of the intentional to the non-intentional in order to make genuine explanatory progress.

In order to begin to see how the possibility of endowment with content requires objectivity, we need to enlist the distinction between semantics and metasemantics and focus on the latter.[5] As it is commonly understood, semantics is concerned with specifying semantic contents and their modes of composition, whereas metasemantics is concerned with the general issue of content-determination. An easy illustration of the distinction is afforded by the semantics and metasemantics of proper names. Direct reference theorists claim against descriptivists that the contents of names are simply their bearers. So the content of "Bismarck" in my mouth, say, is the man Bismarck. Such identification belongs to semantics. It specifies the content of the name as the entity named. But now the metasemantic question arises: How does the name "Bismarck"

---

[5] The distinction is discussed (under a slightly different terminology) in Joseph Almog, "Semantical Anthropology", *Midwest Studies in Philosophy*, 9, 1984, 479-489; David Kaplan, "Afterthoughts", in Almog, Joseph; Perry, John and Wettstein, Howard (eds.), *Themes from Kaplan*, Oxford, Oxford UP, 1989, pp. 565-614, especially at pp. 573-576; Stalnaker, Robert "On Considering a Possible World as Actual", *Proceedings of the Aristotelian Society Supplementary Volume*, 75, 2001, pp. 141-156; Coleman, Jules and Simchen, Ori, "Law"', *Legal Theory*, 9, 2003, pp. 1-41.

in my mouth come to have the man Bismarck as its content? And here the prevalent metasemantic view that accompanies the semantic theory of direct reference is the causal-historical chain view. It holds that the name "Bismarck" in my mouth comes to have Bismarck himself as its content by virtue of a causal-historical chain running back from my current employment of the name, via my own cognitive history, via the sources from whom I acquired the name, be they a history teacher or an author of a book I have read, and their cognitive histories in turn, and further back via sources of sources of sources all the way down to some initial act of naming Bismarck "Bismarck". This causal-historical chain view is a metasemantic thesis, a thesis that must be distinguished from the semantic thesis that the content of the name "Bismarck" is simply the man Bismarck.[6]

The new theory of reference, which is the general framework of metasemantic productivism that I employ, initiated an externalist revolution in our thinking about the cognitive relations between the mind and the world, specifically the relation between the contents of terms and what they are about, namely, their extensions.[7] Traditionally, the relation between contents and extensions was thought to be a species of satisfaction in the formal sense. On this view, the content of a term poses a mere condition that specifies what the term is about by way of satisfaction of the condition.[8] Such contents were thought to be immediately accessible to the mind of the agent, whereas the portions of the world that terms are about were thought to be cognitively once removed. In this way, cognition was thought of as inevitably mediated by conditions entertained in the mind. It was a crucial feature of this outlook that contents do not depend for what they are on what the terms are about, or even on whether they are about anything at all, just as a mere condition can be the condition that it is whether or not anything satisfies it.

---

[6] The most influential statement of the metasemantic thesis regarding proper names is found in Kripke, Saul A., *Naming and Necessity*, Cambridge, Harvard UP, 1980, pp. 90-97.

[7] I aim to remain as neutral as possible on how exactly to think about what contents are. This semantic issue, while a crucial ingredient in any overall metasemantic story concerning content-determination, lies outside the scope of my immediate concerns.

[8] By "mere condition" I mean to rule out *de re* conditions, conditions that depend for what they are on the objects that satisfy them, such as the condition expressed by "identical with O" where the name "O" is understood to be contributing its bearer directly.

On an externalist view of the relation between the mind and the world, aboutness is no longer thought to depend on satisfaction of conditions. The basic idea is that a term is about whatever it is about by virtue of being *of* it. To get an intuitive handle on what this *ofness* amounts to, it is useful to consider the aboutness of photographs.[9] Consider the case of a photograph taken of one of two identical twins. Suppose further that the twins are so similar (or the photograph so imprecise) that had a photograph been taken of the other twin under suitable conditions, it would have been molecule-for-molecule identical to the actual photograph. Thus, as a mere visual condition, the photograph does not discriminate between the two twins. Yet for all that, it is only about one of them. We simply do not think of the aboutness of photographs as a matter of satisfaction of visual conditions. Rather, we think of this aboutness as having to do with the photograph's ofness. The photograph is about the twin it happens to be of. It is about whichever of the two twins was the relevant causal-historical antecedent to the photograph's formation as consequent. One thing suggested by such cases is that we do not in general individuate photographs in abstraction from what they are about. A photograph of the other twin might have been molecule-for-molecule identical to the actual photograph, yet it would still be a different photograph by virtue of being of —and thus about— someone else.

Semantic externalism considers the ofness of expressions to be essential to their aboutness. Whatever their content is ultimately held to be, the aboutness of terms is achieved via their ofness. This means, among other things, that we have a reversal of the traditional view of the relation between content and extension. Traditionally, contents were thought to specify extensions as mere conditions that do not depend for what they are on what the terms are about. But on the externalist outlook, content crucially depends for what it is on what the term is about– so much so, in fact, that many versions of semantic externalism simply identify contents with extensions. To summarize the contrast in a word we might say that whereas the traditional view thought of content as determining extension, the new orthodoxy thinks of extension as determining content.

---

[9] Such heuristic appeal to photography is inspired by a similar appeal made in David Kaplan, "Quantifying In", *Synthese*, 19, 1968, pp. 178-214.

With these cursory remarks on the general framework in hand, we can now turn our attention to the metasemantics of general terms such as common nouns and adjectives. How are we to think of content-determination for a typical general term? To the extent that we think that those contents depend for what they are on what the terms are about, *i. e.* on their extensions, extension-determination is going to play a crucial role in the overall account of content-determination. Take the familiar example of "water". To the extent that the content of "water" depends for what it is on its extension, an account of how "water" comes to apply to all and only samples of water will occupy a central role in any plausible story about the way in which "water" gains its content. For the remainder of this paper, my earlier antireductionist remarks must be borne in mind. Specifically, the explanatory burden of a metasemantic account of general terms should not be thought of as the reduction of the intentional to the non-intentional.[10]

Elsewhere Jules Coleman and I have defended an account of extension-determination that takes its main cue from Putnam's work in *The Meaning of* "*Meaning*".[11] As against the traditional view that knowing the content of a typical general term is a matter of knowing an extension-fixing criterion that all and only samples of the relevant kind satisfy, Putnam (and, independently, Kripke) has argued convincingly that there is little reason to think that proficient speakers are in possession of any such criteria. For example, adult speakers of English who are proficient with the noun "gold" seldom know a general criterion that applies to all and only instances of gold. But for all that, "gold" applies to all and only instances of gold. How is this determination achieved? It is achieved in two ways: socially and environmentally.

---

[10]   One reason for being pessimistic about the prospects of a naturalistic reduction of aboutness is that for many terms content-determination proceeds by way of linguistic deference to a relevant expertise, as we shall see. Such deference implicates an elaborate authority structure, and there is good reason for being pessimistic about the explanatory prospects of attempting to account for such social phenomena naturalistically, within the vocabulary of cognitive science, say. Such pessimism carries over to a general pessimism about a cognitive-scientific reduction of aboutness. This is an area of heated controversy that obviously demands far more attention than I can devote to it here.

[11]   See Coleman, Jules and Simchen, Ori, "Law", Putnam's classic paper is collected in Putnam, Hilary, *Mind, Language and Reality: Philosophical Papers*, vol. 2, Cambridge, Cambridge UP, 1975, pp. 215 and 227.

Socially: Let us grant that proficiency with "gold" does not entail possession of an extension-fixing criterion. Yet I am a proficient speaker who would be quite easily taken in by samples of iron pyrite ("fool's gold"). Does it not follow that 'gold' in my mouth picks out anything that I would be inclined to regard as gold, including samples of fool's gold? – Not at all. "Gold" in my mouth still applies to all and only samples of gold because what "gold" in my mouth applies to is not just a matter of how things are with me considered in isolation from the rest of my linguistic community. Rather, it is an intricate matter of social exchange that Putnam dubbed "division of linguistic labor" and Gareth Evans likened to the relation between producers and consumers. I will follow received practice and refer to this phenomenon as "linguistic deference". The basic idea is that ordinary speakers ("novices") successfully refer to gold and not to fool's gold by employing "gold" via their tacit reliance on a relevant expertise, in this case metallurgy. By placing their trust in an expert doctrine, speakers can employ general terms to refer determinately to things regarding which they are relatively ignorant. Reference cannot be such a difficult cognitive task so as to require each and every member of the linguistic community to become an expert on what is being talked about.

Environmentally: Had speakers' surroundings been relevantly different, say with some distinct yet superficially indiscernible water-like substance occupying the role of water, or with cat-like demons occupying the role of cats, the extensions of "water" and "cat", and so their contents, would have been different from what they actually are. Speakers employ such terms to speak about whatever in the world is around them. We employ "cat" to refer to cats and not to cat-like demons. Our counterparts in the cat-like demon world use "cat" to refer to cat-like demons and not to cats. The difference is in what is around. Moreover, if we bear in mind that our cats are not demon-cats and consider our own intuitions regarding whether or not "cat" as spoken by us applies also to demon-cats, assuming such things are possible, the answer is a resounding No, *even under the further assumption that we would never be able to tell them apart from cats*. "Cat" as spoken by us applies to cats and to nothing else. Similar intuitions can be elicited for other general terms. This strongly suggests that there is an indexical element in extension-determination for a typical general term. To fall under the extension of "cat" is to bear some relation —a relation that demon-cats, for example,

do not bear— to paradigmatic cats in the environment that are referred to indexically, say by employing 'this[the furry meowing thing over here]'. In other words, the actual interaction between speakers and their environment determines what 'cat' applies to in any possible world.[12] As Frank Jackson puts it regarding the example of 'water': "The reference in all worlds is settled by what is watery and the subject of the relevant acquaintance in the actual world" (39).[13]

Putting the above two points together yields the following schema of an account of extension-determination for general terms. The extension of a typical general term "N" is specified by the following condition:

(†) $N(x) \leftrightarrow (x, this)$.[14]

"N" applies to all and only those items, in any possible world, that bear a relevant similarity relation ˜ to whatever is indexically referred to by "this" as spoken in the actual world. Take "gold". It applies to all and only samples of a substance, in any possible world, that bear a certain similarity relation —in this case microstructural similarity— to paradigmatic samples of the substance indexically referred to by "this" as spoken in the actual world. In short, "gold" refers to whatever is microstructurally close enough to paradigmatic samples of gold in speakers' actual environment. And discerning that the relevant similar-

---

[12] (For those interested in semantic scruples): When considering whether "cat" applies to occupants of other possible worlds that are relevantly similar to actual cats picked out indexically, or whether instead it applies to occupants of other possible worlds that are relevantly similar to items picked out indexically *in those other worlds*, the first option seems to be supported by, and the second option to conflict with, a basic semantic fact about indexicals, namely, that indexicals take large scope relative to intensional operators. Thus, (i) is consonant with the logic of indexicals as it is commonly construed, whereas (ii) is not (we let the square brackets indicate scope and "˜" stand for the relevant similarity relation):

    (i) [this]N("x)(cat(x) « ˜(x,this))

    (ii)N("x)(cat(x) « [this] ˜ (x,this))

    If this is correct, then the rigidity of 'cat' depends on a feature of its metasemantics, a feature that depends, in turn, on a feature of the semantics of indexicals.

[13] Frank Jackson, *From Metaphysics to Ethics*, Oxford, Clarendon Press, 1998.

[14] Nothing requires that "this" be the particular indexical expression involved in extension-determination for a typical general term. The choice of a specific indexical is only for heuristic purposes, as should become clear from the discussion to follow of the role of (†) within the overall metasemantic account.

ity relation (call it "$_M$", for "metal") indeed obtains is something that is left to metallurgical expertise to decide. In other words, the social aspect of extension-determination for "gold" enters primarily[15] in discerning that $_M$ obtains between putative instances of gold and paradigmatic samples, whereas the environmental aspect enters in the employment of the indexical "this" in the actual presence of paradigm instances of gold.

Before moving on to consider further wrinkles, I need to say more about the role of (†) itself within the overall account of content-determination on offer. First and foremost, (†) is not meant itself to capture the content of "N". It is offered as an extension-fixer that speakers are tacitly committed to, given their actual referential intentions. (†) figures in a metasemantic, rather than a semantic, account of extension-determination. The idea is that speakers employ a typical general term "N" *as if* they are committed to (†) as an extension-fixing stipulation. And the force of this "as if" claim is just that speakers employ "N" with the referential intention to pick out whatever is in fact relevantly similar to paradigmatic instances of the kind in their environment.

Now, (†), all by itself, is completely schematic. Precisely how we are to think of the implicated referential intentions in concrete instances is a subtle matter, but this much is relatively clear. Referential intention attribution is a species of intention attribution more generally, which is, in turn, a species of attitude attribution. It is a commonplace in attributing attitudes to agents that we attribute to them only those attitudes that they can be expected to have given their overall epistemic situation. What is rather surprising and seldom noticed is that this constraint, mild as it may seem, actually renders certain common philosophical attributions of referential intentions highly implausible. One such implausible attribution is the attribution to ordinary speakers of the metaphysical realist intention to employ "water" to refer to anything relevantly similar to paradigmatic instances of water from the standpoint of the world as it is in itself, beyond whatever we might come to believe about the matter. Another implausible attribution, from the other end of the philosophical spectrum, is the attribution to ordinary speakers of the radical subjectiv-

---

15 It is sometimes suggested that linguistic deference enters not only in discerning relevant similarity to paradigmatic instances of the kind, but also in identifying the paradigmatic instances themselves. This seems correct for some deferential terms and incorrect for others. The general issue need not concern us here.

ist intention to employ "water" to refer to whatever the agent herself would regard as relevantly similar to paradigmatic instances of water in a way that is not susceptible to any external check on the matter. (I will come back to radical subjectivist referential intentions below). Neither attribution is plausible in light of speakers' other attitudes. Referential intention attribution is an exercise in making sense of speakers' attitudes within an overall metasemantic story about how a given term comes to possess the content that it has. More specifically, it consists in squaring referential intentions with speakers' other intentions, beliefs, desires, hopes, fears, etcetera. It should thus aspire to remain as true as possible to speakers actual attitudes and should refrain as much as possible from subjugating them to extrinsic philosophical agendas.

A related point concerning the role of (†) is that (†) is neutral as to whether "N" is linguistically deferential or not. Certain terms, such as the natural kind term "gold", are linguistically deferential if any term is. As mentioned above, individual proficiency with "gold" does not require of speakers to be capable of discerning that a given sample of substance is relevantly similar to paradigmatic instances of gold. Rather, speakers are best understood to be implicitly deferential to metallurgy to decide on such matters. But other terms do not exhibit linguistic deference – the non-natural kind term "chair" is one salient example. It is no part of our linguistic practices *vis-à-vis* "chair" that we are deferential to some chair-expertise to discern relevant similarity to paradigmatic chairs. In "Law" Jules Coleman and I argue that despite initial appearances to the contrary, the deferential-non-deferential distinction cuts across the natural-non-natural distinction. It is just not the case that the natural kind terms are the deferential ones whereas the non-natural kind terms are the non-deferential ones. Some natural kind terms, such as "puddle", are not deferential, whereas some non-natural kind terms, such as 'carburetor', are.[16] This makes the question of what determines whether or not a given term is linguistically deferential more demanding than is often presumed. Much of the discussion in "Law" is devoted to answering this question and drawing implications from the answer for the recent con-

___
[16] I use "natural kind term" in a way that purports to remain neutral with respect to further metaphysical commitments regarding natural kinds. Puddles are not non-natural; hence, they are natural. (I realize that there is a lot more that can be said here concerning what in general counts as natural, but it is not required for present purposes).

tention in the philosophy of law that we are as linguistically deferential to jurisprudential expertise with respect to 'law' as we are to metallurgy with respect to "gold".[17]

To summarize this all-too-brief metasemantic sketch, extension-determination for a typical general term "N", which plays a crucial role in its content-determination, is achieved via referential intentions to pick out whatever is relevantly similar to (*i. e.* bears to) paradigmatic instances of N in speakers' environment. (The role of "this" in (†) is simply to make the environmental aspect of extension-determination salient). If 'N' is linguistically deferential, speakers leave it up to some relevant expertise to discern whether or not obtains. If "N" is not linguistically deferential, speakers are relatively self-reliant in this regard. Either way, "N" gains its extension, and consequently its content, via referential intentions that specify that it is to apply to anything bearing ˜ to paradigmatic instances of N in speakers' environment.

We are now finally in a position to explore the bearing of this metasemantic story on the question of objectivity. If the above sketch is on the right tracks, then content-determination depends on extension-determination. And extension-determination depends, I now claim, on an objective measure of similarity to paradigmatic instances of the relevant kind. In other words, it is built into the possibility that the general term "N" should have whatever content it happens to have that there be an objective measure of similarity to instances of the kind. Extension-determination for a typical general term depends on the existence of some independent standard that can facilitate a genuine difference between cases where instances only seem to be relevantly similar to one another and cases where this is in fact the case. To see this, we turn to consider some examples.

In the metasemantic literature it is often presumed that in the case of substance terms such as "water" or "gold", such a standard is provided by the microstructure of the substance. Take "gold" again. It is thought to apply to all and only samples that are microstructurally close enough to paradigmatic samples of gold. But this is a matter that can and does easily transcend mere seeming similarity to gold. Ordinary proficient speakers are not privy to the procedures whereby experts distinguish samples

---

[17] For a defense of this contention see Nicos Stavropoulos, *Objectivity in Law*, Oxford, Clarendon Press, 1996.

of genuine gold from samples that only seem to the unaided mind to be samples of gold. It is precisely here that linguistic deference enters the picture. For it is built into our linguistic practices vis-à-vis these substance terms that in determining whether or not such a term applies in a given case, an ordinary speaker is beholden to a relevant expertise on the matter. This means that distinguishing similarity to paradigm instances in such cases from mere seeming similarity is a matter that novices leave to experts to decide.

The case of simple artifact terms such as chair' or "hammer" is more difficult, but here, too, there are means for distinguishing similarity to paradigm instances of the kind from mere seeming similarity. As mentioned above, a term such as "chair" is not linguistically deferential, so such means are not provided by an expert doctrine. In addition, while as a linguistic community we are obviously successful in classifying things under the label "chair", how we determine relevant similarity to paradigmatic chairs is not entirely transparent to us. However, there are compelling empirical reasons for thinking that relevant similarity in such cases is heavily informed by the intended function of instances of the kind.[18] Let us assume that this is correct: A proficient speaker employs "chair" with the referential intention to pick out anything that is relevantly similar to paradigm instances, where relevant similarity has much to do with being intended to serve the same function as paradigmatic chairs. But whether or not a given item is intended to serve the same function as a paradigmatic chair is a matter that can easily transcend what merely seems to be the case. If it is indeed true, as findings on conceptual development suggest, that relevant similarity to chairs is a matter that is heavily informed by intended function, then in the case of an item that despite appearances to the contrary has no intended function, speakers would stand corrected if they initially classified it as a chair and were then informed that in fact the item has no intended function.

Be that as it may, whether we are dealing with linguistically deferential kind terms or with ones that are not, without facilitating a distinction between genuine similarity to instances of the kind and mere seeming similarity, no extension could be secured for the kind term in question,

---

[18] The psychological literature on conceptual development abounds with attempts to identify features of artifacts that are generally considered to be essential to them. See, for example, Keil, F. C., *Concepts, Kinds, and Cognitive Development*, Cambridge, MIT Press, 1989.

and, consequently, no content. For without any means of effecting a
seems-is distinction in extension-determination, a term would apply to
anything seeming to be relevantly similar to (what seem to be) paradig-
matic instances of the kind. This means that no possibility of erroneous
application of the term would be facilitated. But without the possibility
of genuine error in application, there is no possibility of genuine correct-
ness in application either. In this way, the term would not contribute to
the truth-conditions of claims in which it partakes; that is to say, it would
lack content altogether.

This line of reasoning bears repetition in a more orderly and sche-
matic fashion: Let "N" be a putative kind term for some putative kind
N. Suppose that (i) there is no objective measure of similarity among in-
stances of N. Then, (ii) there is no objective measure for membership in
N. Thus, (iii) there is no possibility of genuine error in the application of
"N". But then, (iv) there is no possibility of genuine correctness in the
application of 'N' either. Therefore, (v) "N" does not contribute to
the truth-conditions of claims in which it partakes. And therefore, (vi)
"N" lacks content.

A standard subjectivist response to this argument is to claim that noth-
ing in it effectively rules out the possibility of endowment with subjec-
tive content. In other words, so the objection presses on, nothing that
has been said so far rules out that a general term may have such content
that is shaped by the radical subjectivist referential intention alluded to
above, the intention to pick out anything that merely *seems* to be rele-
vantly similar to what merely *seem* to be paradigm instances of the kind.
All that is needed for endowment with subjective content is that the rele-
vant similarity relation itself be subjective. For "N" to be endowed with
subjective content, its extension need only be fixed by the subjective in-
clination to regard things as relevantly similar to whatever one is subjec-
tively inclined to regard as paradigmatic instances of seeming-N. As
long as this remains a standing possibility, it is just wrong to claim, as I
have, that endowment with content requires an objective measure of sim-
ilarity in extension-determination.

Let us examine what would transpire if we withdrew the requirement
that there be an objective measure of similarity in extension-determina-
tion, in the way suggested by the objection. Let us assume for the sake
of argument that the attribution of the radical subjectivist referential
intention is in fact adequate for some substance term that applies to all

and only instances of *seeming water-like*. Let this term be "water*". We are supposing, then, that "water*" applies to whatever is deemed by the speaker to be relevantly similar to (*i. e.* whatever bears * to) other instances of water*, where being an instance of water* is constituted by merely seeming water-like.

Now, in order for any general term "N" to gain a determinate extension, and so to contribute to the truth-conditions of claims in which it partakes, there has to be some way of effecting a seems-is distinction that would allow us to say that something can only *seem* to be relevantly similar to instances of N without actually being relevantly similar to those instances. We just saw how a seems-is distinction is provided for in the examples considered above of substance terms and simple artifact terms. In the case of a term such as 'water*', what would be required is a facility to distinguish cases of merely seeming to bear* to instances of water* from cases of genuinely bearing* to them. But here comes the crucial point: If something seems to bear* to instances of water*, then *ipso facto* it bears* to them! For to bear* to something is to seem relevantly similar to it. But to seem to bear* to something is to seem to seem relevantly similar to it. But seeming to seem relevantly similar to something is just to seem relevantly similar to it all over again. Seeming does not genuinely iterate. Consider any case of seeming to f. If something seems that it seems to f, then it also thereby seems to f. If something seems that it seems red, then it also seems red; if something seems that it seems sweet, then it also seems sweet; if something seems that it seems painful, then it also seems painful. In other words, there is no place to insert the requisite seems-is wedge when it comes to seeming-to-f. If I believe that something is red then I may be mistaken, for it may only seem to me that it is red while being some other color. But if something only seems red to me – where "seems red" is not just a stylistic variation on "is red" – I cannot be mistaken about that. What this means, in effect, is that extension-determination for "water*", and so content-determination, cannot take place after all. So the metasemantic question regarding "water*" remains unanswered. If this is correct, then despite initial appearances to the contrary, a term such as "water*" can-

not gain content. The possibility of endowment with content requires objectivity in extension-determination. [19]

I began this paper by briefly considering the oppositional way that writers on objectivity typically set things up, using their strategy as a point of contrast to my alternative metasemantic strategy. Rather than focus on judgments in a particular domain and raise the question of the metaphysical status of the facts that constitute their subject matter, followed perhaps by the question of epistemic access to those facts, I chose to focus on the world-thinker interplay that undergirds endowment with semantic content. Given the framework of metasemantic productivism sketched above, it turned out that the very possibility of such endowment requires that there be some independent measure to facilitate the difference between genuine similarity to instances of a given kind and mere seeming similarity. In this way, objectivity is required for extension-determination, and so required for content-determination more generally. But curious minds still want to know: What *kind* of objectivity do legal facts, let us say, enjoy, as opposed to the objectivity of the facts of natural science, or of moral facts, or of mathematical facts? Is the objectivity of legal facts not "softer" in some sense than the objectivity of natural-scientific facts? Yet for all that such questions may strike us as gripping and unavoidable, it is far from obvious that anything useful can be said about objectivity as a feature of facts considered as truth-makers for our judgments. We can, however, turn our attention to objectivity as presupposed by the possibility of endowment with content of specific terms and perhaps learn thereby something important about their employment.

---

[19] The transition from the claim that if "water*" seems to apply to something then it does in fact apply to it, to the claim that "water*" has no content, might give rise to the following worry. Suppose that "water*" applies by seeming to apply. Could I not still *misapply* it, say by intending to misapply it? But in that case, it seems that a genuine contrast between application and misapplication for "water*" can be facilitated, in which case "water*" can gain a determinate extension, and so a determinate content, after all. However, further reflection will reveal this to be gratuitous. Under the conditions specified above, what might it mean to say that we can misapply "water*"? Suppose I resolve to misapply it in a given instance. In what (or against what) might my misapplication of it consist? The only available answer is that "water*" seems to misapply in this given instance. In other words, the term in question applies by seeming to apply and misapplies by seeming to misapply. And this can only mean that there is no talk of genuine application or misapplication here. (Thanks to Mark Greenberg for drawing my attention to this worry).

By way of conclusion, let us turn again to "negligence". If the overall argument of this paper is sound, and to the extent that the metasemantic framework sketched above is independently plausible, then in order for "negligence" to gain its content there must be an independent measure to distinguish genuine similarity to paradigmatic instances of negligence from mere seeming similarity. The presumption that 'negligence' is endowed with content, that it contributes in relevant ways to the truth-conditions of claims in which it partakes, entails, among other things, that "negligence" can apply wrongly. As long as speakers —and this includes judges— consider the term "negligence" in their mouth to have content, as long as they consider it as contributing in relevant ways to the truth-conditions of claims in which it partakes, they must also consider it to be susceptible to misapplication in specific cases. If the term is not susceptible to misapplication, it cannot have content. And it is, in fact, overwhelmingly likely that speakers, including judges, consider such terms to be susceptible to misapplication. In this, "negligence" is not so different from "water". As long as speakers —and that includes chemical experts— consider the term "water" in their mouth to have content, as long as they consider it as contributing in relevant ways to the truth-conditions of claims in which it partakes, they must also consider it to be susceptible to misapplication in specific cases. If the term is not thus susceptible, then it cannot have content after all.

Where does this all leave us? One thing that can be said about objectivity is that a measure of independence from what we happen to deem relevantly similar to paradigmatic instances of kinds is presupposed by the very possibility that our kind terms are endowed with content. And this includes what we deem relevantly similar to paradigmatic instances of kinds even at some hypothetical end of inquiry, or under ideal epistemic conditions. The epistemic conditions under which the verdict that Jones was negligent was reached may have been ideal. Yet for all that, in order for the term 'negligence' in the court's mouth to have content, that is, to contribute to the truth-conditions of claims in which it partakes, the court must be regarded as susceptible to error in application. This means that despite all the epistemic ideality in the world that may happen to obtain as regarding the court's employment of "negligence", it is compulsory to treat it as answerable to some independent standard that can facilitate the distinction between being relevantly similar to paradigmatic instances of negligence and merely seeming to be so.

Does this conclusion saddle us with some version of untenable Platonism? – Hardly. For the metasemantic strategy for broaching objectivity does not begin with the customary examination of the metaphysical underpinnings of the relevant facts, followed by a treatment of epistemic access to them. The above argument, if indeed successful, illustrates that objectivity is presupposed by the possibility of endowment with content. To my mind, given the attractiveness of the metasemantic picture offered above, such a conclusion is on much firmer ground than are customary defenses of objectivity, defenses that tend to raise far more perplexities than they actually succeed in quelling.