

In Patrick Greenough and Michael P. Lynch, eds., *Truth and Realism* (Oxford, 2006), 137-161.

Abundant Truth in an Austere World

Terry Horgan and Matjaž Potrč

University of Arizona, University of Ljubljana

What is real? Less than you might think. We advocate *austere metaphysical realism*—a form of metaphysical realism claiming that a correct ontological theory will repudiate numerous putative entities and properties that are posited in everyday thought and discourse, and also will even repudiate numerous putative objects and properties that are posited by well confirmed scientific theories. We have lately defended a specific version of austere metaphysical realism which asserts that there is really only one concrete particular, viz., the entire cosmos (see Horgan and Potrč (2000, 2002), Potrč (2003)). But there are various potential versions of the *generic* position we are here calling austere metaphysical realism; and it is the generic view that constitutes the ontological part of the overall approach to realism and truth that we will describe here.

What is true? More than you might think, given our austere metaphysical realism. We maintain that truth is *semantically correct affirmability, under contextually operative semantic standards*. We also maintain that most of the time, the contextually operative semantic standards work in such a way that semantic correctness (i.e., truth) is a matter of *indirect* correspondence rather than *direct* correspondence between thought or language on the one hand, and the world on the other.¹ When correspondence is indirect rather than direct, a given statement (or thought) can be true even if the correct ontology does not include items answering to all the *referential commitments* (as we will here call them) of the statement.² This means that even if a putative object is repudiated by a correct ontological theory, ordinary statements that are putatively about that object may still be true. For instance, the statement “The University of St. Andrews is in Scotland” can be semantically correct (i.e., true) even if the right ontology does not include any entity answering to the referring term ‘The University of St. Andrews’, or any entity

answering to the referring term ‘Scotland’. This general approach to truth is what we call *contextual semantics*.³

We will here call our package-deal position, comprising austere ontology and contextual semantics, *austere indirect-correspondence realism* (for short, *AIC realism*). In the first section we will elaborate somewhat upon our version of AIC realism, stressing some prima facie advantages of AIC realism over various alternative approaches to truth and realism that are currently on offer in philosophy. The remainder of the paper will be organized around three kinds of skeptical challenge that can be raised against AIC realism. We will seek not only to *meet* these challenges, but also to use them as guides for further developing—and for motivating—the specific version of AIC realism that we seek to defend.

Briefly, the three challenges are these. First, our appeal to contextual variability of semantic standards faces an unhappy dilemma: the need to make a choice between (i) claiming (implausibly) that thoughts and statements that seem categorical (e.g., ‘The University of St. Andrews is in Scotland’) are really implicitly relativistic and non-categorical in content, or instead (ii) claiming (again implausibly) that thought and discourse harbor extensive un-noticed semantic ambiguity (since the specific meaning of one’s thoughts and statements depends so heavily, according to contextual semantics, upon implicit contextual factors). Second, AIC realism, because it advocates austere ontology, threatens to be *grossly* contrary to common-sense beliefs about the world, and/or grossly contrary to what the best science seems to tell us about the world. The AIC realist, after all, denies the existence of vastly many of the entities posited in ordinary thought/talk and in scientific thought/talk. Third, skeptical doubts are apt to arise about whether an adequate general account—in terms of general, systematic, normative principles—can be given of matters like (i) indirect-correspondence semantic standards of various sorts, and (ii) the dynamics of contextual variation in semantic standards.

Concerning the first challenge, we will bring to bear some ideas recently deployed by Horgan and Timmons (2002) in an effort to make sense of Putnam’s notion of conceptual relativity. We will harness these ideas as a proposed way of going between the horns of the dilemma.

Concerning the second challenge, we will have quite a lot to say about how common sense beliefs and scientific beliefs fare, under the approach we advocate. We will be arguing (i) that AIC realism accommodates *most* of common sense and science very well, (ii) that AIC realism plausibly explains why common sense balks so strongly at austere ontology, and that the balking reaction (because it is thus explainable) does not constitute good grounds for rejecting AIC realism, and (iii) that common sense itself, when it turns ontologically reflective, actually generates strong theoretical grounds for an austere ontology.

Concerning the third challenge, we will argue that semantic normativity in most contexts, and likewise the dynamics of contextual variability in semantic normativity, very likely are too subtle and complex to conform to general, systematic, principles. As we will put it, semantic normativity is very likely *quasi-particularistic*, rather than rule-like (cf. Potrč 2000, Potrč and Stahovnik 2004). The case for quasi-particularism about semantic normativity rests partly upon ontologically reflective common sense: certain common-sense reflections that motivate an austere ontology also motivate semantic quasi-particularism. We will also argue that those who maintain that there *must* be general rule-like semantic principles, in order for humans to be capable of mastering semantic normativity, are making highly dubious empirical assumptions about the workings of human cognition.

1. Contextual Semantics, Metaphysical Realism, and Ontological Austerity.

There are various ways to call into question the idea that all referential commitments, in language and thought, are really ontological commitments. There are paraphrase strategies of one sort or another. There are fictional approaches, and error theories. There are epistemic theories that seek to reduce truth to epistemically warranted affirmability, or to some idealized variant of it. There are global irrealist approaches which effectively deny that there's any such thing as genuine ontological commitment at all, understood in a metaphysical realist way. And there are approaches to truth according to which issues about truth and issues about ontological commitment have very little directly to do with one another. Let us briefly say something about various *prima facie* advantages of AIC realism over these various other

approaches. Along the way we will bring into focus some further features of our favored version of AIC realism.

There are various ways of trying to “paraphrase away” discourse that is referentially committed to putative entities that one considers ontologically dubious. One way is to offer paraphrases that just drop out the relevant referring terminology altogether (as in paraphrasing ‘She has a charming smile’ as ‘She smiles charmingly’). Another is to offer paraphrases that effectively identify the erstwhile offending entities with entities one considers ontologically more respectable (as in identifications of numbers with *sets* of one sort or another).

But there are reasons to be very dubious about the paraphrase approach as a general strategy for avoiding questionable-looking ontological commitments. One problem is that often there are no terribly plausible candidates for paraphrasing. How, for instance, might one plausibly paraphrase the statement ‘The University of St. Andrews is in Scotland’ into some statement that eschews university-talk and nation-talk? A further problem is that often, among the *marginally* eligible candidate paraphrases, there will be far too many that look *equally* (albeit marginally) eligible. For instance, there will be too many equally eligible ways to paraphrase talk about the University of St. Andrews into talk about things—or sets or mereological sums of things—like people and buildings and computers and vehicles and such.

An obvious advantage of AIC realism, in comparison to the paraphrase approach, is that the former eliminates the *need* for systematic “paraphrasing away” of discourse with referential commitments that one has reason to think are not genuine ontological commitments. Instead of reformulating the relevant claims in an ontologically austere discourse, one instead goes “soft on truth” for the relevant discourse: one claims that the discourse operates under *indirect-correspondence* semantic standards (IC semantic standards)—and that *truth*, for a discourse governed by such standards, is just semantic correctness *under* those standards. The upshot is truth without ontological commitment, and without the need for systematic paraphrasing.

Lately there has been some enthusiasm for “semantic pretense” theories of ontologically questionable thought and discourse: theories that treat such thought/discourse as being effectively a form

of *fiction* (e.g., Walton 1990, Crimmins 1998). We take it that according to these views, numerous beliefs that we common-sensically hold true are not *really* true at all—just as it’s not really true that there’s a person named Santa Claus who lives at the North Pole and dispenses presents at Christmas, or that there was a person named Sherlock Holmes who lived on Baker Street in London and was a brilliant sleuth. An obvious *prima facie* advantage of AIC realism, in comparison to semantic pretense theories, is that the former allows us to respect the persistent belief that numerous ordinary beliefs are *literally true*, whereas semantic pretense theories do not. AIC realism has the same advantage over *error* theories of common-sense thought and talk.

Epistemic reductionist theories of truth seem to exert a perennial attraction for some philosophers. Recently influential versions include Putnam’s some-time contention that truth is “ideal warranted assertibility” (Putnam 1981), and Wright’s some-time suggestion that truth is identical to a form of idealized warranted assertibility that he calls “superassertibility” (Wright 1987). But such approaches face at least two very serious objections, not faced by AIC realism. First, unless one idealizes all the way up to something like a “God’s-eye epistemic vantage point,” there will be persistent cases where truth and idealized warranted assertibility evidently diverge—for instance, (i) statements about the distant past for which no extant evidence exists one way or the other, (ii) statements about certain goings-on in distant portions of spacetime outside the light-cone of the human race, etc. Second, even if one idealizes so much that idealized warranted affirmability coincides—or necessarily coincides—with truth (given the extent of idealization), it will nevertheless be the case that when a statement is ideally warrantedly affirmable, this will be *because* it is true—rather than its being true *because* it is ideally warrantedly affirmable. (This is what Wright (1992) calls “the Euthyphro contrast.”) As we say, AIC realism does not face these problems—which constitutes a very considerable advantage of AIC realism over epistemically reductionist accounts of truth.

Here we should pause to stress a certain feature of contextual semantics, as we construe it: semantic standards of correct affirmability are likely to be *closely intertwined* with epistemic standards of warranted affirmability, even though the former are not *reducible* to the latter. Such intertwining is

entirely to be expected: since systematic true belief is an evaluative ideal with respect to normative epistemological notions like justification and warrant, semantic standards and epistemic standards ought to “fit” one another. Warranted thoughts/statements will be ones that are *likely to be true given the available evidence*—i.e., likely to be *semantically correct* given the available evidence. Small wonder, then, that contextually operative epistemic standards of *warranted* affirmability will be closely intertwined with contextually operative semantic standards of *correct* affirmability—notwithstanding the fact that one can’t *reduce* semantically correct affirmability to epistemically warranted affirmability (or to some idealization of the latter).⁴

Another, related, point to stress is this: Facts about contextually operative standards of epistemic warrant can be expected to be a good guide to contextually operative semantic standards (given the fit between them). Consider, for instance, what would need doing in order to obtain good epistemic warrant for the claim that Warner Brothers Films is owned by the Miramax Entertainment Corporation. The relevant evidential standards just *don’t* require obtaining good evidence that there are entities, viz., WARNER BROTHERS FILMS and MIRAMAX ENTERTAINMENT CORPORATION, included among the furniture of the universe. (Here and occasionally below, we resort to Putnam’s capitalization convention, as a way of signaling *ontological* uses of terms that, according to our general view, often are employed without ontological commitment because the contextually operative semantic standards governing these terms can be indirect-correspondence standards.)

One way to stop worrying about which referential commitments of thought and language are genuine ontological commitments is to claim that *none* of them are, and that the idea of mind-independent, language-independent, world is a metaphysical extravagance. This is global metaphysical irrealism (which is often combined with epistemically reductionist views of truth). But we ourselves find global irrealism itself to be a doctrine so metaphysically extravagant as to be not only wildly implausible but also well nigh unintelligible. Since these claims have been argued for elsewhere (Horgan 1991, 2001b), we will not repeat the arguments here.

Lately there has been much enthusiasm for one or another version of “minimalism” or “deflationism” about truth—roughly and generically, the view that the various instances of the Tarskian schemas T and F pretty much exhaust all there is to the concepts of truth and falsity (e.g., Field 2001, Horwich 2001). One major disadvantage of minimalism/deflationism is this: although there certainly are schema-T uses of the truth predicate, there are important correspondence uses too. For instance, for someone who holds that moral judgments/statements have nondescriptive overall content, *one* important right thing to say about them with respect to truth—under *one* legitimate usage of ‘true’—is that *they are neither true nor false*. A significant advantage of AIC realism over minimalism/deflationism—more specifically, an important advantage of the *contextual semantics* component of AIC realism over minimalism/deflationism—is that contextual semantics can smoothly accommodate *both* schema-T uses of the truth predicate and correspondence uses, for modes of discourse for which (according to nondescriptivist treatments of these modes of discourse) these uses diverge.

Here we will mention a few aspects of how this accommodation works.⁵ Horgan and Timmons distinguish between *tight* and *non-tight* contextual semantic standards. A judgment/statement is governed by *tight* semantic standards if those standards conspire with how the world is to render the statement/judgment semantically correct or semantically incorrect; otherwise, the semantic standards are *non-tight*. (If, for example, humor judgments are partially expressive of the judge’s own sense of humor and thus can vary among several people without anyone’s being mistaken, then the semantic standards governing thought and talk about what’s funny are non-tight.) Second, contextually variable semantic standards govern the truth predicate itself. Third, under one usage of ‘true’ that is sometimes contextually appropriate—a *correspondence* usage—a statement/judgment whose governing semantic standards are non-tight counts as neither true nor false. (Horgan and Timmons claim that under this usage, *moral* statements/judgments are neither true nor false in their standard usage, since normally they are governed by non-tight semantic standards.) Fourth, on *another* usage of the truth predicate that is sometimes contextually appropriate, the truth predicate conforms to schema T. (Horgan and Timmons claim that on this usage, truth ascriptions to moral judgments/statements are *morally engaged* meta-level

judgments/statements, as are the first-order judgments/statements to which they are ascribed. Such a truth predication is a fusion of semantic and moral evaluation.)

Vagueness also makes trouble for standard versions of deflationism/minimalism about truth. It is extremely plausible that thought-contents and statements applying a vague category to a “borderline case” are neither true nor false. Suppose, for instance, that Jones is borderline-bald (and thus is also borderline non-bald). Then natural thing to say about truth and falsity is that the statement ‘Jones is bald’ is *not true*, and also is *not false*. But the deflationist/minimalist is hard pressed to avoid reasoning in the following way:

Suppose that ‘Jones is bald’ is not true. By schema T, ‘Jones is bald’ is true iff Jones is bald.

Hence Jones is not bald. By schema F, ‘Jones is bald’ is false iff Jones is not bald. Hence, ‘Jones is bald’ is false.

But the conclusion of this reasoning seems wrong. Since Jones is a borderline case of baldness, the statement ‘Jones is bald’ is neither true *nor* false.

In sum: We have briefly set forth some significant *prima facie* advantages of AIC realism, in comparison to a range of other approaches to truth and realism now on offer. Doing so has involved some degree of further elaboration of the position beyond what we said at the outset. In particular, we have stressed that under contextual semantics, semantic and epistemic normative standards are apt to be closely intertwined, even though semantic standards are not reducible to epistemic ones (or to some idealization thereof). We have also stressed the tight/non-tight distinction, and the way it allows for contextually variable semantic standards to govern the truth predicate itself. In particular, given a non-descriptivist treatment of moral discourse, it allows for both correspondence uses of the truth predicate and schema-T uses.

2. Contextual Semantic Variation: Content Relativity vs. Identity-Preserving Difference.

We turn next to some remarks about how to think about the workings of contextually variable semantic standards, drawing upon material from Horgan and Timmons (2002). One idea that commonly

gets invoked in connection with context dependence is that contextual factors are an implicit aspect of the very *content* of a thought or sentence involving such factors—so that the actual content is really *relativistic* content. For example, when one says that the table is flat, one really means that it is flat relative to such-and-such standards. When one says that Miramax Entertainment Corporation owns Warner Brothers Films, one is saying that *relative to such-and-such semantic standards*, Miramax Entertainment Corporation owns Warner Brothers Films.

This approach seems just wrong. The trouble is that even when contextually variable semantic parameters are in play, typically the discourse/thought governed by such parameters operates *categorically*: one is speaking/judging from within a semantically committed stance in which one *accepts* the semantic standards as operative. When one says that something holds relative to such-and-such standards, on the other hand (even if the ‘relative to such-and-such standards’ part is implicit), one need not be accepting those standards at all. People do sometimes talk/think in a relativistic way, but when they do so they typically employ *explicit* relativization.⁶

But if indeed discourse/thought governed by contextual parameters is normally categorical in content rather than implicitly relativistic, then the following question naturally arises: What is the semantic relation between a statement/thought that is affirmed under such-and-such contextually operative semantic standards, and one (expressed the same way) that instead is affirmed under so-and-so alternative semantic standards? A looming worry is that this contextualist approach is effectively positing massive *equivocation* in the use of concepts and terms that are governed by contextual parameters. Commonality of content, across different uses, is threatened.

Our position on this matter is as follows. Inter-context variation in operative semantic parameters (variation in aspects of what Lewis (1979) called the “score in the language game”) constitutes *identity preserving* change/difference in meaning and concepts. This is what Horgan and Timmons (2002) call *differance* in meanings/concepts (adapting Derrida’s terminology), rather than what’s ordinarily called ‘difference in meaning/concept’. The idea is to think of meanings and concepts (whatever exactly they are) as items that can exhibit, from one context of usage to another, certain *identity-preserving*

differences—much as other kinds of entities (e.g., persons) can differ in certain ways over time (e.g., in hair color).

To illustrate such difference consider the following two-part remark that might be made by philosopher while discussing ontology and common-sense belief: “Are there corporations? *Of course!* Are there *really* corporations? No!”⁷ One can assert both parts of this remark in close succession to one another, but not in the same breath. There is a subtle semantic variation at work across the two uses of ‘corporation’, but not a semantic difference that constitutes anything nearly as great as ordinary semantic ambiguity (as with ‘bank’ as financial institution or edge of a river). The word ‘corporation’ has the *same* meaning when employed under different semantic standards governing the first part of the remark and the second part respectively, even though this contextual variation is indeed a difference: it’s an identity-preserving difference in meaning. Likewise, *mutatis mutandis*, concerning the *concept* expressed by ‘corporation’.

Another way one might put the point is this. Certain terms, as used in two different token statements made under somewhat different settings of the contextually variable score in the language game, are *weakly synonymous* but not *strongly synonymous*. One can also use this terminology about the respective thought-contents expressed by the statements. Statements/thoughts that are weakly but not strongly synonymous do exhibit certain meaning/content differences, but these (again) are much less stark than ordinary semantic ambiguity—and the differing uses exhibit much more semantic *commonality* despite these differences.

3. Accommodating Common Sense, Mostly.

Let us now consider the status of common-sense thought and talk, within a picture that weds contextual semantics to an austere ontology. First, we will argue that common sense gets *very largely* reconciled with austere ontology, under our proposed contextual semantics. Second, we will argue that contextual semantics plausibly explains why and how common sense balks at the claims of austere ontology—an explanation that accommodates something importantly *right* about thus balking, while also

revealing the subtle ways that common sense confuses itself in thus balking (and in addressing ontological issues). Third, we will argue that there are powerful considerations from within common sense itself that generate serious theoretical pressure toward some kind of austere ontology—i.e., an ontology that eschews numerous entities to which common sense itself is referentially committed. We will take up the first of these themes in this section, and the next two in sections 4 and 5 respectively.

It is important to realize that AIC realism delivers full fledged, bona fide, *truth* despite the austerity of the ontology. Truth under indirect-correspondence semantic standards (IC semantic standards) is not some sort of second-class semantic status, something lesser than full fledged or genuine truth. To think otherwise is to be too much in the grip of the very picture of truth that is here being repudiated. Indirect correspondence is perfectly fine, as a mode of correspondence. There are numerous ways that people engage in referential commitments that do not actually involve *ontological* commitment, and doing so often serves people’s thinking/communicating/living purposes very well. Think, for instance, of talk of universities, nations, corporations, religions, musical works of art. We “quantify over” such entities freely in various modes of thought/discourse, and in general the contextually operative semantic standards certainly do require of the world that it be one way rather than another, in order for a thought/statement involving such referential commitments to count as semantically correct—i.e., as true. But in order for the world to be one of the ways it might be that would render true our ordinary thoughts/claims about such entities (in typical contexts of usage), it simply is not required that the right *ontology* should include entities answering to terms like ‘St. Andrews University’, ‘Scotland’, ‘General Motors Corporation’, ‘Catholicism’, or ‘Mozart’s 27th piano concerto’.

How then is one to think of truth, qua indirect correspondence? How is one to understand the contention that the world makes the statements true, but not in a way that involves ITEMS answering to the statements’ referential commitments? There is actually a familiar conception of truth conditions that can here be invoked. Think of *truth conditions* for a thought/statement as constituted by something like a set of *ways the world might be* (“possible worlds”) such that if the world is one of *those* ways then the statement/thought is true (false).

There are various ways one might “do the ontology” of such a conception of truth conditions. For example, each of these ways the world could be may get construed as a *maximal cosmos-instantiable property*. One would probably want to refine such a picture in various ways, we suspect—for instance, to accommodate vagueness and indexicality. (So-called “centered” ways the world might be, with a designated location to be thought of as the location of the thinker/speaker, could perhaps accommodate indexicality. One might need to soften up on the idea of a truth condition as a determinate *set* of ways the world might be, to accommodate vagueness.) But the core idea remains: A statement *S* with various referential commitments in it has a truth condition *TC(S)*—roughly, a set of maximal cosmos-instantiable properties—such that *S* is true provided that the cosmos instantiates one of these properties in *TC(S)*, and is false provided that the cosmos does not instantiate one of these properties in *TC(S)*. The truth conditions for a statement *S* with referential commitments to universities, nations, family trusts, or the like etc. can perfectly well be *satisfied*—i.e., the cosmos can perfectly well instantiate one of the properties in the set *TC(S)*—even if the furniture of the cosmos does not include such putative entities as UNIVERSITIES, NATIONS, or CORPORATIONS.

4. Common Sense vis-a-vis Austere Ontology: The Role of Scorekeeping Confusions.

Common sense balks, and balks strongly, when the claims of some austere ontology are put forward. If we say to you, “There aren’t really such things as tables, chairs, or people,” you are apt to react by judging such a claim preposterous—not just false, but screamingly false. That this reaction is apt to occur, and to occur with this kind of strength and intensity, is a datum—one that needs to be dealt with, within the overall position we are advocating.

Our position does not directly *accommodate* the datum, in a straightforward manner; i.e., our position does not treat the balking reaction as correct, and does not vindicate it. Rather, on our view you are making a *mistake* if you reject our claim that there are no such things as tables, chairs, or people. For, as we intend the claim (*viz.*, as a claim that is being made under direct-correspondence semantic standards (DC standards)), what we are saying is *true* if indeed some austere ontology is the correct one.

We take it that one serious obligation of a philosophical theory is to give a suitably *respectful* treatment of intuitively plausible judgments/statements that the theory rejects—all the more so if the judgments are as intuitively plausible as are judgments like “*Of course* there really are tables, chairs, and people” (thought/said in response to our own claim that there *are not* really such things). Respectfulness means, inter alia, that although the judgments/statements in question get treated as embodying some kind of error, it’s not a silly or stupid or merely careless error. Rather, the account will be better if the error is treated as particularly subtle, and one that people are particularly apt to make given the normal workings of their cognitive apparatus.

Let us introduce the idea of a *competence-based performance error*. The thought is that although the cognitive system falls short of its own competence in generating the given judgment or intuition or belief, nonetheless it is producing this performance error in a way that in some sense emanates from competence.⁸ Think, for instance, of the Mueller-Lyer illusion. The cognitive system is working the way it is evolutionarily designed to, in producing a perceptual presentation of the one horizontal line as longer than the other. If you form the perceptual belief that one of the horizontal lines is longer than the other, then you are committing a competence-based performance error. Your visual system is working the way it has evolved to work (although it happens here to be yielding up a visual illusion), and you are forming a mistaken perceptual belief by reliance on the presentational content of your visual experience.

There’s something right and something wrong, in your belief-forming process. The something right is that the lines really do *look* to be not the same length. Someone not knowing better would justifiably form the *belief* that they are not the same length. Something has gone wrong, though. The lines have in fact the same length, as can be determined by ordinary measurement techniques. You are competent to tell whether they are the same length or not; they are the same length; and yet here you are, *judging* them to be different in length.

It is important, too, that the lines are apt still to look the different lengths even after the measurement has been carried out. They will still look that way even when you realize that they are

actually the same length. The appearance will persist (or at least there will be a strong tendency for it to do so), even after you come to believe that it is non-veridical.

Return now to the fact that common sense balks, seriously balks, at the claims of austere ontologies. We propose to treat this as also a competence-based performance error, of a certain sort. Although it is *mistaken* to repudiate austere ontologies on common-sense grounds, in a very important sense this is a mistake that is rooted in semantic/conceptual competence itself.⁹

We will describe the basic picture in eight steps. First: As competent users of concepts and terms that are governed by various implicit contextual parameters (aspects of what Lewis called the score in the language game), people typically deal with these parameters so smoothly and competently that they do not even notice them. Witness, for example, Lewis's charming little spiel about Bruce the cat, in the section of Lewis (1979) in which he is illustrating how well we accommodate to the current contextually relevant score that determines which contextually eligible referent of a definite description 'the F' (e.g., 'the cat') counts as most *salient* F (and hence as the referent of 'the F'). He writes:

Imagine yourself with me as I write these words. In the room is a cat, Bruce, who has been making himself very salient by dashing madly about. He is the only cat in the room, or in sight, or in earshot. I start to speak to you:

The cat is in the carton. The cat will never meet our other cat, because our other cat lives in New Zealand. Our New Zealand cat lives with the Cresswells. And there he'll stay, cause Miriam would be sad if the cat went away.

At first, "the cat" denotes Bruce, he being the most salient cat for reasons having nothing to do with the conversation. If I want to talk about Albert, our New Zealand cat, I have to say "our other cat" or "our new Zealand cat." But as I talk more and more about Albert, and not any more about Bruce, I raise Albert's salience by conversational means. Finally, in the last sentence of my monologue, I am in a position to say "the cat" and thereby denote not Bruce but rather the newly-more-salient Albert. (Lewis 1979, p. 241)

We accommodate to score-changes without necessarily realizing that we are doing so, and without necessarily noticing that contextual parameters are in play at all. Witness, for other examples, the other kinds of phenomena Lewis canvasses in that landmark paper, all of which he persuasively argues are ones in which implicit, contextually variable, parameters of “score in the language game” are in play. Humans appear to be hugely good at scorekeeping in most ordinary contexts—so good that they do it without noticing that implicit semantic parameters are in play, and that they vary.

Second: Normally when one employs certain categories, and certain terms expressing them, there is a very strong default presumption in play that the contextually operative semantic standards are ones under which certain affirmative, referentially committing, thoughts/statements *employing* those categories will be true (i.e., semantically correct).

Third: On the view we advocate, contexts in which DC semantic standards are being employed are *extremely* rare and rarefied: they are contexts of serious ontological inquiry.

Fourth: When we say to you “There are no tables, chairs, or people,” there is a strong tendency for you to react to this statement as though it were being made under relatively *typical* contextually operative standards for categories like ‘table’, ‘chair’, and ‘person’. This is partly because of the default presumption just mentioned (the second point), and partly because of the extremely rarefied nature of thought/discourse that’s conducted under DC semantic standards (the third point). There is a strong tendency, that is, for you *not* to accommodate properly to the contextually operative score in the language game that governs our own remark (even though *normally* people accommodate to one another’s remarks smoothly and naturally), and for you instead to respond to us with a remark that itself is governed by more *ordinary* semantic standards (IC standards) for categories like ‘table’, ‘chair’, and ‘person’.

Fifth: So there is something importantly right in your balking reaction. Your claim “Of course there are tables, chairs, and people!” is certainly true, *under the contextually operative standards governing your own usage*. And your claim is also in line with the strong default presumption that a discourse in which familiar, commonly used categories are in play is governed by semantic standards in

which certain affirmative claims employing those categories can turn out to be true (i.e., semantically correct).

Sixth: But even though your balking reaction is in these ways *competence-based*, there is something importantly wrong in it nonetheless. To wit: you have failed to accommodate to the semantic standards that are contextually appropriate, given that serious ontological inquiry is the game we mean to be playing.

Seventh: Also, you fail to appreciate that implicitly contextually variable semantic parameters are in play at all, and that our remark was made under a different language-game score than your own. Here, the fact that you normally accommodate to score-changes so smoothly and naturally that you do not even *notice* score-parameters and their variation is contributing to the *mistake* you are making.

Eighth: The sense of oddness and screaming falsity of claims like “There are no people” is apt to persist, even for someone who accepts (as we do) an ontology that does not include people. (More below on *why* we do.) The various factors cited above explain why this should be so.

This kind of mistake is what we will call a *scorekeeping confusion*. Given the overall picture we are presenting, such scorekeeping confusions are entirely to be expected. This is because of the ways they emanate so naturally from one’s own semantic/conceptual competence, including one’s competence at *normally* handling implicit contextual parameters so smoothly and naturally that one typically does not even notice that they are there.

So the upshot is that your balking reaction is being respectfully explained away as a certain kind of competence-based performance error—a scorekeeping confusion on your part. Although *this* aspect of common-sense belief gets repudiated by our account, rather than being vindicated, our account has the significant virtue of explaining nonetheless why this aspect of common-sense thinking occurs—and also why it will likely continue to exert a strong intuitive pull, even upon those who come to accept that the right ontology is some kind of austere ontology.

The larger upshot so far, concerning common sense, is (1) that common sense is *mostly* accommodated within the picture we are pushing (as argued in the preceding section), and (2) that

common sense's mistaken tendency to balk at austere ontological claims can be plausibly explained as a subtle kind of competence-based performance error—viz., a scorekeeping error.

5. Common-Sense Pressures toward Austere Ontology.

Let us now argue that common sense itself, when it goes reflective about matters ontological, generates strong pressure in favor of an austere ontology. This can and does happen even without the realization that semantic standards are contextually variable, or the realization that contextually operative semantic standards can require either direct correspondence or some form of indirect correspondence, or the realization that most contexts of thought and discourse are ones in which indirect-correspondence standards prevail. In sections 5.1 and 5.2 we will be incorporating material from Horgan and Potrč (2000) lightly edited to play up the common-sensical sources of the considerations being advanced.

5.1. *Metaphysically Lightweight Posits*

We will use the phrase 'metaphysically lightweight posit' in a deliberately vague way. Under this rubric we include "socially constructed" institutional entities like corporations, universities, nations, and multi-national organizations (e.g., NATO). We also include various non-concrete cultural artifacts, like Beethoven's fifth symphony (as distinct from concrete performances of it) and Quine's book *Word and Object* (as distinct from concrete tokens of it).

It is not plausible to common sense that institutional entities like corporations and universities are denizens of the world-itself, *over and above* entities like persons, buildings, land masses, items of office equipment, and the like. Yet, when one considers whether it might be possible to "reduce" a putative entity like a university to these other kinds of entities—say, by identifying each university with some set of them (or some "mereological sum" of them), or by systematically paraphrasing statements that posit universities into statements that do not—there is no plausible reductive account remotely in sight. For, the project of systematically paraphrasing university-talk into statements that eschew all talk of universities looks hopeless; and the trouble with attempts to *identify* a university with some set (or sum) of buildings, persons, computers, etc. is that there are always numerous equally eligible candidate-sets (or candidate-

sums), and there is no reason to identify the university with any one of these over and against any of the others. Likewise, *mutatis mutandis*, for other kinds of institutional entities like corporations and nations, and for non-concrete cultural artifacts like Beethoven's fifth symphony and Quine's *Word and Object*.

So the appropriate common-sensical conclusion to draw, about thought and discourse employing these metaphysically lightweight posits, is that the posits do not pick out items in the world-itself. Common sense does not, however, thereby embrace the idea that thoughts and statements about such entities don't *correspond* to the world, don't depend on how the world is. On the contrary: there are certain ways the world might be such that, if it's one of these ways, then the statement 'General Motors is a multinational corporation' is *true*; and there are other ways it might be that would render this statement *false*. It's just that the correspondence-constituting ways the world might be need not include such items as CORPORATIONS or NATIONS. Common sense viscerally realizes this—which means that common sense viscerally appreciates that the truth of the statement about General Motors is a matter of *indirect* correspondence.

Of course, one could accept these common-sensical-looking conclusions while remaining a robust ontological realist about middle-sized dry goods like tables and chairs, about persons and other living organisms, and about the posits of physics and the special sciences. But there is more to come.

5.2. *The Non-Arbitrariness of Composition*

Peter van Inwagen (1990) wields an important and powerful form of metaphysical argumentation that has been too little appreciated in philosophy. He poses what he calls the Special Composition Question (for short, the SCQ): "When do several objects jointly compose an object?"

Van Inwagen considers several initially plausible candidate-answers to the SCQ. He argues that each has highly implausible consequences, viz., commitments to putative entities that are not genuine objects at all according to our usual ways of thinking and talking. (For example, the suggestion that *contact* among a group of objects is what makes them jointly compose an object entails the grossly

counterintuitive result that when two people shake hands, a new compound object comes into existence that ceases to exist when their hands separate.)

Van Inwagen argues, by elimination-via-counterexample of various initially plausible potential answers to the SCQ, that the only acceptable answer is that several objects compose an object when they jointly *constitute a life*. On this basis, he concludes that the right ontology of physical objects includes only two kinds of material beings: (1) “simples,” whatever these might turn out to be (e.g., electrons and quarks, perhaps), and (2) living organisms. Discourse that posits other kinds of concrete objects, he says, should be understood by analogy with talk about the motion of the sun through the sky: useful and informative, but not literally true.

Two important theoretical desiderata are in play, in van Inwagen’s discussion of the ontology of material beings: (1) finding a systematic, general, answer to the SCQ, and (2) adopting an ontology that conforms reasonably well to our pre-theoretic common-sense beliefs, and our scientifically informed beliefs, about what kinds of physical objects there are. Van Inwagen argues very persuasively that these desiderata are deeply in tension: they cannot both be satisfied. But he also assumes, without explicit argument, that insofar as the two desiderata conflict, satisfying (1) is theoretically more important than satisfying (2).

Now, it might well be asked whether one *should* attach more theoretical importance to obtaining a general and systematic answer to the SCQ than to “saving” tables, chairs, and other objects that we pre-theoretically consider robustly, mind-independently, real. If so, why? We submit that the answer to the first question is affirmative, and emerges from common sense itself when common sense goes reflective about metaphysics. Here is what common sense realizes: An adequate metaphysical theory, like an adequate scientific theory, should be systematic and general, and should keep to a minimum the unexplained facts that it posits. In particular, a good metaphysical or scientific theory should avoid positing a plethora of quite specific, disconnected, *sui generis*, compositional facts. Such facts would be ontological surds; they would be metaphysically queer. Even though explanation presumably must bottom out somewhere, it should bottom out with the kinds of “unexplained explainers” we expect to find in

physics—viz., highly general, highly systematic, theoretical laws. It is just not credible—or even intelligible—that explanation would bottom out with specific compositional facts which themselves are utterly unexplainable and which do not conform to any systematic general principles. Rather, if one bunch of physical simples compose a genuine physical object, but another bunch of simples do not compose any genuine object, then there must be some reason *why*; it couldn't be that these two facts are themselves at the explanatory bedrock of being.

There cannot, then, be a body of specific compositional facts that are collectively disconnected and unsystematic, and are individually unexplainable. Such ontological arbitrariness is not possible in the world-itself—the world whose constituents are van Inwagen's concern. In Horgan (1993) this is called *the principle of the non-arbitrariness of composition*. This principle is fundamental and highly plausible, and is a very compelling general requirement on theory construction. It generates the requirement that an adequate metaphysics of concrete particulars be one for which there is a general and systematic answer to the special composition question. This requirement has very strong weight in metaphysical theory-construction, enough to trump the desideratum of preserving the posits of common sense and science. *Common sense viscerally appreciates this*—even though common sense will also balk if someone boldly and overtly *affirms* some austere ontological claim, such as Van Inwagen's claim that the only real physical beings are physical simples and living organisms.

So the upshot so far is that common sense's appreciation of the need to provide a systematic and general answer to the SCQ, together with the great difficulty of providing such an answer that also pretty much includes all and only the kinds of physical objects that are posited in ordinary discourse and in scientific discourse, generates a strong common-sense basis for holding that right ontology of concrete particulars will have to be one that posits either many fewer, or else many more, kinds of concrete particulars than those that are usually posited in science and in common sense. One candidate ontology that provides a systematic and general answer to the SCQ is van Inwagen's own, comprising only physical simples on the one hand, and living organisms on the other.¹⁰

5.3. *The Benign Logical Incoherence of Vagueness*

Elsewhere (e.g., Horgan 1998 and other papers of Horgan's cited therein, Horgan 2002b), it is argued at some length that vagueness is logically incoherent in a certain specific way; that this kind of logical incoherence is *benign* in that it can be quarantined in language and thought in order to avoid problems like becoming committed to rampant contradictions; but that the logical incoherence inherent to vagueness renders *ontological* vagueness—i.e., vague OBJECTS and/or vague PROPERTIES—impossible. Here we will summarize the key line of thought very briefly.¹¹

Essential to vagueness are the following features. First, *sufficiently small differences don't make a difference*. They don't make a difference, for instance, with respect to the applicability of a vague category (e.g., 'bald', or 'part of Mount Everest') to items that differ from one another in sufficiently small ways (e.g., two people one of whom has one more hair on his head than the other, or two pebbles on the ground one of which is three inches further from the peak of Mount Everest than the other). Second, *sufficiently large differences do make a difference*. Vague categories do discriminate among cases that are sufficiently different from one another—for instance, between bald folks (who have little or no hair on their heads) and non-bald folks (who have plenty), or between pebbles that are part of Mount Everest and far-away pebbles that are not. But third, *iterated combinations of sufficiently small differences add up to sufficiently large differences*—as illustrated by the kinds of small-difference iterations that generate instances of the sorites paradox. When common sense reflects clear-headedly on these three principles, common sense realizes (i) that all three principles are indeed essential to genuine vagueness, (ii) that the three principles jointly embody a form of logical incoherence, and hence (iii) that no genuine OBJECT or PROPERTY could possibly be vague. (Epistemicists reject claim (i) and maintain instead that vague categories always delimit perfectly sharp, albeit unknowable, boundaries—a view *wildly* contrary to common sense.) Common sense also appreciates, at least viscerally, that people nonetheless get by well enough with vagueness in *thought-content* and in *language*, provided that they are suitably judicious in employing vague concepts and terms—which includes systematic avoidance of sorites-paradoxical reasoning.

But as common sense also realizes, most of the referential commitments of common sense itself—and numerous of the referential commitments of scientific theory as well—are to *vague* items. Concreta like tables, chairs, persons, stars, and cells, for instance, all are vague in certain ways—e.g., with respect to their temporal boundaries, and with respect to their composition. Categories like baldness, tallness, redness, and numerous scientific categories are vague with respect to which concreta do and do not fall under them. So common sense appreciates, at least viscerally, that the right ontology *cannot* include such items, and thus *must* be an austere ontology.

Direct-correspondence semantic standards surely cannot tolerate logical incoherence. When one is thinking/talking under such standards, sorites arguments become sound *reductio ad absurdum* arguments against the existence of vague objects or properties. This fact thus reveals that in *most* contexts of thought and utterance, the contextually operative semantic standards are *indirect-correspondence* standards of one sort or another—since most contexts are ones in which numerous thoughts and statements employ vague concepts and terms in ways that carry referential commitments.

Although direct-correspondence semantic standards cannot tolerate the kind of logical incoherence endemic to vagueness, we contend that indirect-correspondence standards *can* do so—and can operate in such a way that the incoherence gets “quarantined” rather than leading to malevolent consequences (e.g., rampant commitment to contradictions). There’s lots that needs to be said about such quarantining, and about the workings of indirect-correspondence semantic normativity that incorporates the specific kind of incoherence that’s essential to vagueness. (For instance, such normativity cannot be understood as yielding, in combination with how the world is, some full and determinate distribution of semantic statuses over all truth-evaluable statements and thought-contents; for, any such total distribution, as long as it respects the principle *sufficiently large differences make a difference*, would end up positing sharp status-transitions of one sort or another, regardless of how many semantic statuses are in play.) But pursuing such matters is a task for another occasion.

For present purposes, the upshot of what has been said about vagueness is this. On one hand, common sense itself can reflectively appreciate the kind of logical incoherence that’s the very essence of

vagueness—and thus can also reflectively appreciate that ontological vagueness is a priori impossible. On the other hand, common-sense’s—and science’s—rampant referential commitments to ontologically vague objects and properties are entirely compatible with austere ontologies that repudiate all vague OBJECTS and PROPERTIES—*provided* that austere ontology is wedded to the claim that in most contexts of thought and discourse—including most scientific contexts—the contextually operative semantic standards are indirect-correspondence standards. AIC realism is just such a package-deal ontological/semantical position.

6. Quasi-Particularist Semantic Normativity.

Various questions/demands are apt to be posed to us about indirect-correspondence semantic standards, and also about the dynamics of intra-contextual determination of—and trans-contextual variation in—such standards. One kind of question/demand would be this: “What are the *general principles* that systematize intra-context semantic normativity of any given kind, and what are the *yet more general principles* by which specific normativity is determined *within* any given context and is altered *across* any given change in context?”

Why suppose there are, or need to be, such general principles in order for our position to be viable? One reason that is apt to lie behind the demand for such principles is the idea that without them, semantic normativity simply would not be *learnable* and *masterable* by human beings.

Yet another kind of question/demand might be posed in the following way: “You claim to be offering a picture whereby indirect-correspondence semantic standards relieve thoughts and statements from being *ontologically* committed to objects and properties to which they are *referentially* committed. Well, in order to make good on such a picture, you need not only (i) to articulate general principles, but also (ii) to articulate these general principles *in an austere vocabulary that eschews any referential commitments of its own to items other than those posited by an austere ontology*. For, unless and until you can deliver such *ontologically austere* formulations of the relevant general principles, you will not have

really eliminated ontological commitment to the various putative objects and properties that are eschewed by austere ontologies.”

Those who pose such questions/demands are apt to think that unless and until we provide the answers being asked for, our position will not really *be* a full-fledged position so much as a mere *skeleton* of a position. They are apt to think that the skeleton needs some meat on its bones before it will be worthy of serious consideration, and before it can even be assessed for intelligibility.

Also, those who pose these questions/demands might be skeptical that principles of the kind being asked for even *exist*; and on that basis, they might doubt that there really is any conceptually stable, viable, position about truth and ontology of the kind we have been seeking to delineate here.

Our own view (cf. Potrč 2000, Horgan 2002b, Potrč and Strahovnik 2004) is that (i) the kinds of demands we have mentioned probably cannot be met, but (ii) they very probably *don't need* to be met. We doubt very much that semantic normativity conforms to fully general principles *at all*—let alone to fully general principles that could be stated in an ontologically austere vocabulary. We also doubt very much that semantic normativity *needs* to conform to such principles in order to be learnable and masterable by humans.

One powerful-looking line of argument for the claim that semantic normativity does not *in fact* conform to exceptionless general principles can be extracted from Van Inwagen's investigation of the Special Composition Question (the SCQ). If there are exceptionless general semantic principles that are conformed to by the kind of semantic normativity normally at work in common-sense thought and discourse, then these principles presumably ought to generate some kind of exceptionless general answer to the SCQ—an exceptionless general answer that nicely systematizes people's various common-sense judgments about whether or not, in various specific cases, a bunch of things jointly compose a whole thing. But the apparent lesson of van Inwagen's careful investigation of the SCQ is that *common-sense judgments about composition do not conform to any such general principle*. (That is why one gets driven to an austere ontology, insofar as one seeks a systematic general answer to the SCQ.) Rather, common sense judgments about object composition are something of an unsystematic hodgepodge—which

presumably means that the semantic-normative standards these common sense judgments reflect are *also* an unsystematic hodgepodge—which presumably means, in turn, that as far as object-composition judgments/statements in ordinary contexts is concerned, *semantic normativity does not conform to exceptionless general principles*.

The point can be generalized. For thought and discourse in general, there is just no obvious reason why the semantic normativity that governs it, in ordinary contexts, needs to be fully systematizable via exceptionless general principles. On the contrary, we think there is a heavy burden of proof upon those who would insist that it *must* conform to such principles; and we would argue that that burden cannot be discharged.

We will use the expression ‘*quasi-particularist semantic normativity*’ for forms of semantic normativity that cannot be fully systematized by exceptionless general principles. The point of the prefix ‘quasi’ is to allow for the possibility that to some extent the relevant kind of normativity can be *partially* systematized by *non-exceptionless* general principles—principles that are vague enough to allow for exceptions (whether or not these principles have explicit ‘*ceteris paribus*’ clauses in them). Such partially-systematizing general principles might be at work *within* contextually specific forms of semantic normativity, and/or they might partially systematize the dynamics of trans-contextual variation in semantic normativity. A nice example, concerning trans-contextual variation, is Lewis’s *rule of accommodation for presupposition*, which he formulates this way: “If at time *t* something is said that requires presupposition *P* to be acceptable, and if *P* is not presupposed just before *t*, then—*ceteris paribus* and within certain limits—presupposition *P* comes into existence at *t*.”¹²

Why think that quasi-particularist semantic normativity is not what prevails, in the case of normal human thought and discourse? Why think that there *must* be fully exceptionless general principles at work? One reason that is apt to come to mind—briefly mentioned already—is the thought that without the latter kinds of principles, the relevant sort of normativity would be too complex and unsystematic and idiosyncratic to be learnable and masterable by humans. We will call this the *learnability argument*.

We think there is strong reason to believe that the learnability argument is seriously mistaken. Its premise is false: learnable/masterable semantic normativity need not conform to exceptionless general principles. And its conclusion is false too: the kinds of semantic normativity normally at work in human thought and discourse just *don't* conform to such principles; rather, the operative kinds of normativity are more subtle, more complex, and more nuanced than that.

Why do we think so? Let us briefly sketch our reasons. There has been a persistent pattern of failure in computational cognitive science to generate plausible models of what Jerry Fodor (1983) calls “central” cognitive processes—processes like (i) the rational generation of new beliefs on the basis of prior background information plus newly acquired sensory information, and (ii) rational planning. As Fodor has persuasively argued for quite a long time—e.g., in the late parts of Fodor (1983) and more recently in Fodor (2001)—and as was argued at some length (in elaboration of Fodor’s argument) by Horgan and Tienson (1996), there is good reason to believe that the problems with attempts to computationally model such central cognitive problems are *in-principle* problems. These problems involve the fact that the relevant kind of information-processing normally needs to be highly *holistic* in nature, potentially drawing upon virtually any item of information the cognitive system might possess, and sometimes (e.g., in making comparative-simplicity assessments of particular hypotheses) drawing upon highly holistic features of *large bodies* of information (e.g., large bodies of information relative to which the given hypotheses are effectively being assessed for relative simplicity).

Horgan and Tienson argue in their book that the problem lies with the assumption that the relevant sort of holistically information-sensitive processing is *computation* over representations—that is, the assumption that such processing conforms to *exceptionless rules* of symbol-manipulation, rules expressible as a computer program. They also describe a non-classical framework for cognitive science, drawing upon certain ideas from connectionism and from the mathematics that goes naturally with it—viz., dynamical systems theory. (This framework retains the idea of a “language of thought,” although with syntax conceived somewhat differently than within the computational paradigm).

Within this framework, central cognitive processes like belief-formation normally will be too complex and subtle (especially in the ways they accommodate holistic aspects of background information) to conform to exceptionless general rules. Cognitively competent cognitive-state transitions will be only *partially* systematizable, via certain inherently exception-ridden generalizations with built-in *ceteris paribus* clauses—generalizations we call *soft laws*. Such soft laws are not—and cannot be refined into—exceptionless algorithmic rules.

Tying this back to various forms of normativity that humans can and do learn and master (epistemic normativity, moral normativity, semantic normativity): the very same kinds of reasons that tell against the idea that competent human cognition is computation over mental representations—is a matter of conformity to exceptionless general rules of symbol manipulation—are also reasons to deny that the competent human deployment of *normativity* (moral, epistemic, semantic,...) is a matter of conformity to exceptionless general rules. Thus, the kinds of reasons that favor a non-computational approach to human central cognitive processing also favor a particularistic or quasi-particularistic approach to the kinds of normativity to which humans are subject—and in particular, to *semantic* normativity.¹³

Let us elaborate a bit. To a very great extent, the same kinds of holistically information-sensitive cognitive processing that are needed for intelligent belief-formation and intelligent planning are also needed in the intelligent deployment of concepts and terms under contextually appropriate semantic standards, and in the intelligent (if implicit) appreciation of which semantic standards are themselves appropriate in various contexts of thought/language one might find oneself in. For, to a very great extent, semantic appropriateness *itself* is apt to be particularistically—or anyway, quasi-particularistically—dependent on holistic elements of the cognitive agent's total situation.

Let us elaborate the point still further, and from a somewhat different perspective: As we see it, the fundamental claim of particularism and quasi-particularism about a given type of normativity is this: *supervenience* relations linking the non-normative to the relevant kind of normativity are not systematizable in terms of compact, general, exceptionless, cognitively surveyable principles. Now, suppose there *are* such supervenience principles, in the case of semantic normativity. Well then, those

very principles should make possible the *paraphrasability* of thoughts/statements about ontologically dubious entities via thoughts/statements about ontologically respectable entities. For, the general supervenience principles should generate, as specific *instantiations*, compact, cognitively surveyable, *truth conditions* for each thought/statement about ontologically dubious entities; *and these very truth conditions would constitute a paraphrase*. (The truth condition for a given statement would be the disjunction of the potential supervenience bases for its truth. And often enough in context, it might be clear enough *which disjunct* is the potentially relevant one, so that this disjunct alone could serve *in context* as the paraphrase.) But in general, as we emphasized early in this paper, acceptable paraphrases for thoughts/statements about ontologically problematic entities just do not seem to be available—and do not *need* to be available, in order for such thoughts and statements to be useful to, and semantically masterable by, humans. The supervenience of semantic normativity upon ways the world might be need not, and very likely does not, conform to short, sweet, exceptionless, general principles. Rather, such supervenience is very likely particularistic or quasi-particularistic in nature.

It is on these grounds that we repudiate the kinds of demands that we mentioned at the outset of this section, and we embrace instead a quasi-particularist version of contextual semantics. Do there need to be general and exceptionless semantical rules that entail, as instantiations of these rules, compactly articulable, cognitively surveyable, and austere expressible *explicit formulations* of the contextually operative truth conditions for any given statement that has referential commitments to entities that are repudiated by an austere ontology? We say, “No, not at all.” What is needed, rather, is just this: (i) the class of ways the world might be needs to be reasonably cleanly partitioned (modulo vagueness) into ways that *conform* with the given statement under the contextually operative normativity, and ways that *fail to conform* with the given statement under the contextually operative normativity, and (ii) a cognitively competent agent adequately good at *judging* such conformity or non-conformity, to the extent that it is accurately indicated by the agent’s own available evidence. (Recall the intimate intertwining of epistemic and semantic normativity.)

So on this approach to contextual semantics, the demands we mentioned at the beginning of this section are illegitimate, as are the presuppositions that motivate them. We admit that we cannot meet those demands, but we deny that this fact constitutes good reason for rejecting our view, or for claiming that the view is a mere skeleton position without adequate meat on its bones, or for denying that we have carved out a genuine and distinctive philosophical position about realism and truth.

7. Conclusion.

When common sense goes reflective, powerful considerations emerge from common sense itself in support of an austere version of metaphysical realism. If such a metaphysical position is wedded to the view that truth is typically indirect correspondence to the world rather than direct correspondence, the resulting package-deal position—austere indirect-correspondence realism—turns out to accommodate ordinary beliefs surprisingly well. The key to such accommodation is this: when the contextually operative standards for semantic correctness require only indirect correspondence to the world rather than direct correspondence, the referential commitments of one's thought and discourse are not ontological commitments.

Demands to spell out general normative principles that govern the intra-context workings of semantic normativity, and general principles that govern the dynamics of trans-contextual variation of semantic standards, rest on dubious presuppositions about both semantic normativity and human cognition. Central cognitive processes appear to be too complex and subtle to conform to programmable general rules over mental representations, and semantic normativity appears to be too complex and subtle to conform to exceptionless general rules.¹⁴

References

- Barnard, B. and Horgan, T. Forthcoming. Truth as Mediated Correspondence.
- Crimmins, M. 1998. Hesperus and Phosphorus: Sense, Pretense, and Reference, Philosophical Review 107, 1048.

- Field, H. 2001. Correspondence Truth, Disquotational Truth, and Deflationism. In Lynch 2001, 483-503.
- Fodor, J. 1983. The Modularity of Mind. MIT Press.
- Fodor, J. 2001. The Mind Doesn't Work That Way: The Scope and Limits of Computational Psychology. MIT Press.
- Graham, G. and Horgan, T. 1994. Southern Fundamentalism and the End of Philosophy (with G. Graham), Philosophical Issues 5, 219-47. Reprinted in M. DePaul and W. Ramsey (eds.), Rethinking Intuition: The Psychology of Intuition and Its Role in Philosophy, Rowman and Littlefield, 1998.
- Horgan, T. 1991. Metaphysical Realism and Psychologistic Semantics, Erkenntnis 34 (1991), 297-322.
- Horgan, T. 1993. On What There Isn't, Philosophy and Phenomenological Research 50 (1993), 693-700.
- Horgan, T. 1998. The Transvaluationist Conception of Vagueness, The Monist 81, 316-33.
- Horgan, T. 2001a. Causal Compatibilism and the Exclusion Problem. Theoria 16, 95-116.
- Horgan, T. 2001b. Contextual Semantics and Metaphysical Realism: Truth as Indirect Correspondence. In M. Lynch 2001, 67-95.
- Horgan, T. 2002a. Themes in My Philosophical Work, Grazer Philosophische Studien 63, 1-26.
- Horgan, T. 2002b. Replies to Papers, Grazer Philosophische Studien 63, 303-41.
- Horgan, T. and Potrč, M. 2000. Blobjectivism and Indirect Correspondence, Facta Philosophica, 2: 249-70.
- Horgan, T. and Potrč, M. 2002. Addressing Questions for Blobjectivism, Facta Philosophica 4, 311-22.
- Horgan, T. and Tienson, J. 1996. Connectionism and the Philosophy of Psychology. MIT Press.
- Horgan, T. and Timmons, M. 2000. Nondescriptivist Cognitivism: Framework for a New Metaethic. Philosophical Papers 29, 121-53.
- Horgan, T. and Timmons, M. 2002. Conceptual Relativity and Metaphysical Realism, Philosophical Issues 12, 74-96.
- Horgan, T. and Timmons, M. In press. Cognitivist Expressivism. In M. Timmons and T. Horgan (eds.), Metaethics after Moore: New Essays in Metaethics. Oxford University Press.

- Horwich, P. 2001. A Defense of Minimalism. In Lynch 2001, 559-577.
- Lewis, D. 1979. Scorekeeping in a Language Game. Journal of Philosophical Logic 8, 339-59. Reprinted in D. Lewis, Philosophical Papers, Vol 1. Oxford University Press, 1983.
- Lewis, D. 1996. Elusive Knowledge, Australasian Journal of Philosophy 74, 549-67.
- Lynch, M. (ed). 2001. The Nature of Truth: Classic and Contemporary Perspectives. MIT Press.
- Putnam, H. 1981. Reason, Truth and History. Cambridge University Press.
- Potrč, M. 2000. Justification Having and Morphological Content. Acta Analytica 24, 151-174.
- Potrč, M. 2003. Blobjectivist Monism. In A. Baechli and K. Petrus, eds., Monism. Ontos Verlag, 125-55.
- Potrč, M. and Strahovnik, V. 2004. Practical Contexts. Ontos Verlag.
- Timmons, M. 1999. Morality without Foundations: A Defense of Ethical Contextualism. Oxford University Press.
- Walton, K. 1990. Mimesis as Make-Believe. Harvard University Press.
- Wright, C. 1987. Realism, Meaning and Truth. Basil Blackwell.
- Wright, C. 1992. Truth and Objectivity. Harvard University Press.
- Van Inwagen, P. 1990. Material Beings. Cornell University Press.

¹ This paper is a broad-brushstroke presentation of the general picture we advocate concerning truth and realism. To some extent we will be summarizing and somewhat re-packaging aspects of the position that has been developing in a series of papers. See, for instance, Horgan (2001b) and several of Horgan's earlier papers cited therein, Horgan and Potrč (2000, 2002), Horgan and Timmons (2002), Barnard and Horgan (forthcoming), Potrč (2003), Potrč and Strahovnik (2004). But we will also incorporate some new ideas, in a way that rests in part on work that Horgan has done with John Tienson on connectionism and the philosophy of psychology, and on ongoing discussions between Horgan and Potrč about particularism, semantic normativity, common-sense realism, and scientific realism.

² We use the expression 'referential commitments' in order to evoke the idea that such commitments arise via the use of logico-grammatical "referential apparatus" in language and thought—e.g., the use in

language of singular terms and existential quantifiers, and the mental use of thought-constituents expressible linguistically by such linguistic items. Read ‘referential’ here as ‘arising from the use of logico-linguistic referential apparatus’—not as ‘purporting to denote some putative item within the correct ontology of the world’. In Barnard and Horgan (forthcoming), the alternative expression ‘ideological commitments’ is employed instead of ‘referential commitments’. (Also, Barnard and Horgan employ ‘mediated correspondence’ instead of ‘indirect correspondence’.)

³ Here we use ‘semantics’ to cover not only language/world relations but also thought/world relations. Semantics concerns matters like *intentionality*, *reference*, and *truth*, both for public language and for thought.

⁴ There is thus intertwining of semantic and epistemic contextual standards. But the contextually operative epistemic standards can vary, to some extent, even when the contextually operative semantic standards don’t. This is because there can be other reasons for contextual variation of epistemic standards. You may be acquainted with the departure time of your airplane. But if your precise arrival at a destination is of extreme importance, you will be taking several additional precautions before you really trust your airplane itinerary. Your epistemic standards will be set much higher. Intuition tells you that semantic standards do not necessarily get higher in such a case.

⁵ For more details, see Timmons 1999, Horgan 2001b, Horgan and Timmons (2000, in press.)

⁶ Indeed, we ourselves talk that way some of the time in setting forth our account of contextual semantics—as when we say, “Truth is semantically correct affirmability *under contextually operative semantic standards*.” Here we should emphasize that we regard the contextual-semantics story as applicable to virtually all the concepts/terms usable to *tell* the story, and in particular to the concept *truth*. Very often, indeed typically, one doesn’t *use* ‘true’ in a relativistic way. On the contrary, one uses it categorically, from within a stance in which one not only accepts certain specific semantic-correctness standards as contextually operative, but one lets them govern one’s truth ascriptions too—the familiar (though not inevitable, on our view) “schema T” usage of ‘true’.

⁷ This is a specific philosophical usage of ‘really’—not the only one that might be appropriate in philosophical contexts, to be sure, and one that is different from *most* uses of ‘really’. This specific usage is employed to overtly signal a shift into a mode of thought/discourse governed by direct-correspondence semantic standards.

⁸ For further discussion of this idea, with applications of it to philosophical puzzlement about the freedom/determinism issue and about the issue of mental causation, see Graham and Horgan 1994 and Horgan 2001a.

⁹ We do not necessarily want to lean heavily on the analogy with the Mueller-Lyer illusion here. That example does put one kind of meat on the bones of the idea of a competence-based performance error. We will be describing *another* kind of meat now, leaving it open how much our account does or does not parallel the Mueller-Lyer case.

¹⁰ A couple of things can be said about van Inwagen’s answer. First, he embraces an austere ontology, as we characterize it early on in this paper. Second, although it does indeed yield a systematic, general answer to the SCQ, we think it runs afoul of what is said about vagueness in the next subsection. It is a vague matter which composites of simples constitute a life, and which do not.

¹¹ We owe the following compact way of putting things to a very helpful conversation with Keith DeRose.

¹² Lewis (1979) treats principles in a more liberal way, as compared to their treatment in Lewis (1996).

¹³ For the question of intertwining between epistemic and semantic normative standards compare footnote 4.

¹⁴ Thanks to Robert Barnard, David Chalmers, and Joseph Tolliver for very helpful comments and discussion.