

The Baltic International Yearbook of
Cognition, Logic and Communication

October 2010 Volume 5: *Meaning, Understanding and Knowledge*
pages 1-83 DOI: 10.4148/biycl.v5i0.278

JODY AZZOUNI
Tufts University,
jodyazzouni@mindspring.com¹

THE RULE-FOLLOWING PARADOX AND THE IMPOSSIBILITY OF PRIVATE RULE-FOLLOWING

ABSTRACT: Kripke's version of Wittgenstein's rule-following paradox has been influential. My concern is with how it—and Wittgenstein's views more generally—have been perceived as undercutting the individualistic picture of mathematical practice: the view that individuals—*Robinson Crusoes*—can, entirely independently of a community, engage in cogent mathematics, and indeed (more generally) have “private languages.” What has been denied is that phrases like “correctly counting” can be applied to such individuals because these normative notions (so the Wittgensteinian analysis is taken to show) can only be applied cogently in a context involving community standards. I attempt to show that this shocking corollary doesn't follow even if Kripke's Wittgensteinian objections to dispositional approaches to rule-following are largely right. My reason for claiming this is that there is another (“sceptical”) solution to the rule-following paradox, one that doesn't favor community standards over individual ones. Furthermore, it doesn't replace truth conditions with assertability conditions; and this latter maneuver is essential to Kripke's sceptical solution favoring the community over the individual.

1. BACKGROUND AND STAGE-SETTING

“Understanding,” as in the “understanding of a concept,” is a word for a complex human capacity that bewilderingly involves subpersonal, conscious/phenomenological, sociological, and normative elements. It's a notion—therefore—that's too dangerously complex to be presupposed in philosophical analyses: it should only be the target of such analyses. Common to several philosophical traditions is the assumption that someone's understanding of a set of concepts is to be grounded in her grasping of the rules that govern those concepts. This captures, according to some philosophers, how we understand the meanings of words, as well as how we understand the concepts involved in simple mathematical practices—such as counting.

Some of those philosophers describe us as having “dispositions” to grasp those rules, or even as having such rules embodied in our dispositions (or as having such rules “embodied as dispositions”) so that we behave appropriately when in situations where we are to exhibit that understanding (e.g., when counting the number of oranges in a basket). A huge literature has focused on “the rule-following” paradox that seems to arise from this picture—that aspects of our rule-following practices seem at odds with the natural picture of our grasping rules, and as a result acquiring dispositions to exhibit those rules in future situations where their application is appropriate. In particular, Kripke's 1982² has generated a large philosophical industry focused on the problems he presented Wittgenstein as raising for any approach that grounds the understanding of simple mathematical rules in the exercise of a set of dispositions.

In the first part of this paper (sections 2-4) I'll describe these problems for dispositional approaches to numerical competence; to a very large extent, the considerations I raise are Kripke's. In so doing, I won't worry about questions of exegesis—of how true to Wittgenstein's concerns Kripke's are.³ Kripke's formulation of the challenges to the dispositionalist is significant regardless of how true it is to Wittgenstein's own thinking on the matter. However, the reader familiar with Kripke's exposition will notice two changes in my presentation of it, one that's very important, and the other less so. First, Kripke's discussion focuses on the central example of a subject adding, and the sceptical challenge that everything she has learnt and previously experienced is consistent

with the possibility that she isn't adding but instead quadding. I substitute for the task of addition the more elementary task of counting, and the corresponding possibility of quounting. This change isn't particularly significant, and indeed, Kripke himself raises the issue in passing, utilizing the term "quounting"—which I've borrowed (see endnote 9).

The second change is rather more significant: my central case focuses on the subject who is counting *things*—objects—and not one who is solving a task involving numerals alone. Children, in fact, learn to count in just this way: they learn to *apply* counting-numbers to the task of recognizing the cardinal number of sets of objects; and it's a kind of knowledge that takes several years to acquire.⁴ One reason that "social solutions" to the rule-following paradox—ones that normatively favor the community over the individual—have been seen by so many philosophers (and sociologists of knowledge) to be the only kind of successful response that can be made to the paradox is precisely because the role of the applications, of the mathematical concepts possessed by most adults, is largely off-stage in Kripke's discussion (and consequently it's largely off-stage in the discussions of the many commentators in this literature).⁵

I should say right off that I'm in partial agreement with one lemma that Kripke draws from the rule-following paradox, and that he himself (78-79) endorses: One must reject "the natural presupposition that meaningful declarative sentences must purport to correspond to facts." I'm in *partial* agreement because I think that rule-following considerations do show that we must accept that *some* meaningful declarative sentences don't correspond to facts; I don't, however, think that they show this with respect to *all* meaningful declarative sentences. (I'll show why the rule-following considerations demand only a *partial* rejection of correspondence in section 15.)

One crucial point is that the dispositionalist hope that so much of Kripke's (1982) analysis is dedicated to dashing locates the *grounds*—of the correspondence relation between meaningful true statements and the facts that they correspond to—within the person's "mind" (broadly described) who is following a rule. That is, the idea is that there are facts *about the person* and what she *can do* that determine the correspondence relations between her meaningful true statements and facts in the world.⁶ Rule-following considerations show (I think) that this

dispositional project fails; dispositions *can't* do the job that most dispositionalists have required of them, and this is largely for the reasons that Kripke (on behalf of Wittgenstein) has given. However, I reject the major thesis that's been drawn from this lemma. Kripke has carefully described himself as only an expositor, but many philosophers (and sociologists of knowledge) have drawn the conclusion that considerations, like the ones I present in sections 2-4, lead inexorably to the result that standards for mathematical practice—and for rule following generally—must be relativized to, or embodied in, the *community* within which individuals learn rules. To use Kripke's (109) language,

What is really denied is what might be called the 'private model' of rule following, that the notion of a person following a given rule is to be analyzed simply in terms of facts about the rule follower and the rule follower alone, without reference to his membership in a wider community.

After I sketch (in section 5) Kripke's exposition of a particular argumentative strategy for this claim—his interpretation of Wittgenstein as replacing truth conditions with assertability conditions—I argue that another solution, another (Kripke, 66) "sceptical solution," is available that doesn't favor community standards over individual standards. I also briefly indicate (in section 4) why a "straight solution" approach—that's meant to yield a result that favors community standards over individual ones—doesn't succeed.

2. SETTING UP THE PROBLEM: THREE CONSTRAINTS ON ANY SOLUTION

Because a subject has apparently learned to count various kinds of collections of various kinds of objects, imagine that we describe her as having learned to count correctly. Thus we take it that she will go on "in the same way" to count in new cases: larger numbers of objects (say) than she has counted in the past, as well as numbers of collections of other kinds of objects (pears instead of oranges, widgets instead of fruit, collections containing assorted heterogeneous items ...).⁷ A component of her acquired counting ability, let's say, is that she has been taught a numeration system that provides a rule for gen-

erating new numerals from earlier ones. Systems of numerals are strikingly different in this respect from the sets of verbal names for numbers that occur in most ordinary languages.⁸ The latter usually belong to a finitary notation for numbers that requires the deliberate coining of new vocabulary for ever-larger numbers, as is the case in English. The possibility of generating ever more numerals, however, is built into numeral systems.⁹ Once a subject has acquired this ability to count, this mastery of a particular numeration system (say), we describe her as *understanding* how to count.

Now the subject attempts to count the items in a collection that's larger than any she's ever counted before: it has 57 items. She gets an answer, the correct answer we would say. Kripke's sceptic challenges the claim that her current answer is in accord with her previous understanding of her task. In the past, Kripke's sceptic claims, she never meant to count, she always meant to quount.¹⁰

To counter this sceptical claim, that the subject intended to quount, not count, and therefore to show that her current answer really *is* in accord with her previous understanding of what she was doing when she undertook the task, what's needed is some grounding fact—or pattern of grounding facts—about the subject that underwrites her previous and current understanding of counting. It's this grounding fact—or pattern of grounding facts—that's to account for it being true in the past that the subject understood how to count rather than it being true that she had acquired the skill and understanding, instead, of “quounting.” This relevant pattern of grounding facts is something about the subject, either something *about* her mental states, or something about the subpersonal events or structures that underlie those mental states, something—in addition—that has changed or developed in such a way as to enable her to now grasp counting. The purported grounding facts about the subject's understanding are supposed to be explained in *these* terms: how the subject's psychological states—what the subject thinks when facing the task—have changed to enable her to understand counting, or how mechanisms or subpersonal representations (neurophysiological or otherwise) have developed or changed or been activated, or whatever, to enable the subject to now be able to count.¹¹

I read Kripke as placing three constraints on the required account of the pattern of grounding facts, dispositional or otherwise, that's to

constitute a subject's understanding of counting (and that's to provide a response to the meaning sceptic). First, there is *the infinitude requirement*. Any subject has counted only finitely many collections of objects. Her presumed ability to count the many collections of objects that she has never counted before raises fresh considerations in two ways. First, as noted earlier, she may have only counted apples and pears, but nothing else: no oranges, and for that matter, no other kinds of objects. Second, she may have counted collections of no more than 57 objects. We take her grasp of counting to include an ability to count collections of objects that differ—in both these ways—from collections she has counted in the past. We would think that something is wrong if a subject hadn't grasped that counting is neutral with respect to the objects she can count: if she felt unable to count objects that were red, or kept in certain boxes, and so on. We would also feel something was wrong if how she counted larger collections of objects deviated from how she counted smaller collections of objects—that (say) she systematically (or *asystematically*) skipped numerals, or started over when she reached a certain numeral, or if she fixated at a particular numeral as the answer, and told us that the rest of the objects weren't relevant.¹² It's part of what we take “understanding” to be, that the subject can handle counting tasks that are different from ones she has handled in the past.

Kripke's second constraint on any account of the grounding facts about the subject that explains how she understands counting is *the justification requirement*. When the subject counts a collection of objects, and recognizes the number of objects that are there, her answer is “justified.” She *ought* to have proceeded in the way that she did given what she *meant*. It's not an accident that her way of counting results in the answer it results in.

Equally crucially, adults and children take themselves to be justified in the answers they give *because* they take themselves to have grasped how to count: they take themselves to understand what “counting” means. If a child (or adult) is asked *why* she counts in the way that she does, she'll say: “Because *that's* how you count.” (Or: “Because *that's* how it's done.” Or: “Because *that's* what counting *is*.”) If a child (or adult) counts in an unusual way, then she'll justify her unusual method by showing (or indicating that she believes) that it always results in

the same answers as strict enumerative counting does.¹³

Notice that this justification that the child (or adult) can offer—that’s what counting *is*—is, except in unusual cases, phenomenologically confident. The subject—once she has learned the procedure—doesn’t hesitate in how she goes about it. And, she doesn’t hesitate in giving an explanation for what she does: This is what counting *is*. Also, we accept this justification from a third-person point of view when we say: She *does* understand what “counting” means. (Or: She *does* understand how to count.)

This confidence in how we count, and in the results we get when we count, is (of course) open to a particular “rider” that everyone recognizes. This is that a “mistake” hasn’t been made. And this brings me to Kripke’s third constraint on any account of the pattern of facts that’s to explain how someone understands how to count: I’ll call this constraint *the mistake requirement*. This third constraint is related to the justification requirement insofar as it may also be described as “normative.” Any subject can make mistakes in counting a collection of objects, and get wrong answers—indeed this is routine. We clearly and naturally distinguish cases where a subject has made a mistake from cases where we describe the subject as not having (fully) grasped the concept of counting (and, notably, we also distinguish cases where a subject has made a mistake—even a systematic mistake—from cases where (i) we would describe the subject as not having yet acquired the concept of counting, and from cases where (ii) we would instead say something seems to be “wrong” *with the subject*). Overlooking an object, or counting one twice (by accident—e.g., because of distraction) are cases we describe as ones where a subject “has made a mistake”; cases where she always deliberately counts all the red objects twice over are cases where something seems more seriously wrong. In the latter sorts of cases, if the subject isn’t just playing games, but thinks she is counting, we are moved to deny that she understands counting. Under certain circumstances, we’ll say she is incapable of learning how to count. An important aspect of our identifying the subject as having made a mistake (although this isn’t a necessary condition) is that the subject can be brought to *recognize* that she has made a mistake. (“You skipped that one.” “Oops,” she says, correcting herself.)

However we describe the subject’s ability to count as embodied in

her, however we characterize the disposition she has acquired by learning to count, room must be made for a concomitant ability to get wrong answers, and to *recognize* they are wrong answers. We don’t normally speak of an “ability” to get wrong answers; but the requirement should be clear nevertheless: We often get the wrong answers even though we grasp counting perfectly well, and any description of the pattern of grounding facts that we use to explain anyone’s ability to count—an ability to execute any task, for that matter—has to allow room for the possibility of errors, and the possibility of recognizing those errors.

I should stress one last related point about how compatible making mistakes—and even systematic *tendencies* to make mistakes—is with our attribution of the understanding of concepts, such as counting, addition, subtraction, multiplication, and division. Some people (as we all know) are amazing calculators: they can quickly and accurately solve numerical problems involving large numbers. It’s not felt that these are people who *grasp* these concepts in a way that the rest of us don’t—more “deeply,” say. It’s just that such people can calculate faster, and know more tricks and shortcuts (in many cases, that they have simply *memorized* more number facts than the rest of us have). Even someone (like myself), who can just about be guaranteed to make an elementary error of some sort when attempting a calculation (any calculation, pretty much) in my head, or even on paper, isn’t regarded as having less of a grasp on the *concepts* of counting, addition, multiplication, and division than calculational *wunderkinds*. What *does* seem to be required, however, to consider someone as possessing these concepts in a way that’s not “defective,” is that she’s open to recognizing the mistakes she’s made. That’s crucial. When someone’s calculations are laid out in detail, and the mistakes are pointed out, if she can’t be brought to *see* that she’s made mistakes, or to understand *why* they’re mistakes, she’s not taken to have the concepts in question—or at least she’s not taken to fully grasp these concepts.¹⁴

3. HOW KRIPKE’S THREE CONSTRAINTS POSE OBSTACLES FOR A DISPOSITIONAL APPROACH

I’ll now illustrate how Kripke’s requirements cause havoc for a natural story about our understanding of counting—that in learning to count,

a subject has acquired a disposition of some kind. It's quite natural to first offer introspective facts as the ones in terms of which the acquired disposition is to be described. The subject is straightforwardly characterized as having recognized the pattern demanded of counting from seeing a finite number of cases; and so she now sees how to go on. But (Kripke, 18), no finite number of cases determines a unique pattern—a unique way to go on. So too, any attempt to describe the subject as *aware* of a rule or of an algorithm that's (i) compatible with the previous cases she has exercised her counting ability upon as well as (ii) determining the answers to future counting tasks she may undertake, fails because any such rule is open to multiple interpretations that cannot be made determinate in terms of how a subject conceives of a rule. Any such rule can be reinterpreted compatibly with what the subject "has in mind," and with all her previous counting experiences.¹⁵

We can see now how Kripke's infinitude requirement rules out a natural description of the needed pattern of grounding facts in terms of what the subject *introspects*. As philosophers have noted, what causes the problem isn't that the subject has to have acquired an ability to provide right answers for *infinitely many* new cases; it's that the subject has to have acquired an ability to provide right answers for *new* cases—ones she hasn't exercised her counting capacity on in the past. The problem is that possible occurrent mental states encompassing the understanding of a rule via the couching of those rules in linguistic forms (even ones involving quantifiers), or via couching rules in some other sort of mental imagery (visual, kinesthetic, etc.), as memories of previous executed tasks, etc., must nevertheless be *applied* to a (new) current task. But how they are *applied* indicates how *they are being interpreted*—and this isn't something determinately fixed ahead of time by what's presently introspectable in the mind.¹⁶

It's at this point that one or another dispositional approach may be reached for. (Part of the motivation for doing so—notice—is to find resources *in* the subject that indicate how the subject is *forced* to the intended interpretation of counting.) I start with a simple construal of such an account of counting:

DIS 1 To grasp 'counting' is to be disposed, when asked for a number of a collection of objects, to give the number of that collection.

This won't do. DIS 1 is compatible with the ability to give the right answers *without* counting. For example, we (humans, infants, and many animals) don't *count* groups of one, two or three items: We *recognize* them.¹⁷ It's certainly imaginable that someone might simply *recognize* indeterminately large groups of items as having certain cardinalities (by the ways—say—that a particular number of objects of so-and-so shapes and sizes must fill out space). Such a person could even have such a skill without realizing—or even being able to realize—that cardinalities of collections can be linearly ordered by magnitude. We must therefore modify the dispositional account to fix on a *particular* way that a subject counts. Here's a candidate:

DIS 2 For someone to grasp 'counting' is for him to be disposed, when asked for the cardinality of a collection of objects, to point at each item in the collection one-by-one and utter a numeral, starting with "1" when pointing at the first item, and to utter the numeral that is the successor of the previous numeral he has uttered each time he points to a new item in the collection, and to give the last numeral he utters when he points at the last item in the collection as the cardinality of that collection.¹⁸

Let's see how DIS 2 fares against Kripke's requirements. First, the infinitude requirement. The problem is that people are very likely *not* to be disposed as DIS 2 requires to determine the number of a collection of objects under all or even most circumstances. If someone is tired, or if the collection is too big, the person may in fact systematically fail to do as DIS 2 requires, or just refuse to try. This may be an actual fact about all people: the dispositions DIS 2 purports to describe are real ones. And, therefore, those dispositions will be sensitive not only to the cardinal sizes of the collections that the subject is faced with, but also to the properties of the items in those collections. For example, if some of the items reflect sunlight too brightly, or are too close together, or are quite tiny, or crawl around on each other, the subject may actually be disposed—not to try to count but—to drop the task altogether (and run for her life). The upshot is that, according to

DIS 2, *no one* grasps counting (because no one is disposed, perhaps in the majority of counting tasks, to do as DIS 2 describes).

As Kripke (30-32) notes, if one introduces counterfactual clauses to handle these cases, one may easily find that DIS 2 gives the wrong answers. It may be that whatever changes *in persons* are needed so that they (counterfactually) attempt counting collections of objects that they are currently disinclined to undertake (e.g., better eyes, differently organized brains, no fear, etc) will instead cause them to deviate from executing the task as DIS 2 requires it to be executed. There is no way we can guarantee that the changes in persons, that shift their tendency to avoid the counting task (in particular cases) to one of engaging in the counting task, don't affect them in other ways with respect to how they undertake that task.¹⁹

Now let's bring in Kripke's third constraint. Indeed, it might be thought that what's (really) wrong with DIS 2 is that it doesn't accommodate, as Kripke's third constraint requires, the empirical fact that we make mistakes. We don't, that is, ordinarily demand of those we take to understand counting (or ourselves, for that matter), that they on any occasion be able to count successfully. We allow exceptions for various reasons including, but not restricted to: the subject tires during the task because of the cardinal size of the collection, the items in the collection have properties that distract the subject from the task, the subject is ill or drunk. We even allow *inexplicable* failures by the subject to complete the task correctly. Despite deviations in these and other cases from the behavior required by DIS 2, and even in one's dispositions to behavior, we are still willing to describe the subject as understanding how to count. Why not, then, simply exclude such cases in "exception-clauses," so that someone is regarded as grasping counting even despite the deviant behavior?

The suggestion begs the question against alternatives. Consider counting* which differs from counting in a number of cases. We can, in exactly the same way, describe the subject as understanding how to count* by taking the cases that the subject tends to deviate from counting* with respect to, and placing *them* in exception clauses. What, then, tells us that people grasp or understand counting instead of it being counting* that they grasp or understand?

The problem with DIS 2 is that it attempts to identify the dispo-

sitions of a subject—when faced with a counting task—with the *correct* procedures for counting. But as we've just seen, the alternative of characterizing the dispositions in question as those that actual subjects have, but supplementing that characterization with exception clauses where such deviate from the correct answers, doesn't fare any better. Defending the claim that the subject understands counting in this way fails because any proponent of the claim that the subject understands some nonstandard form of counting can do exactly the same thing.

Let's turn, lastly, to Kripke's justification requirement. The problem, simply, is that the disposition view offers resources that aren't the ones that anyone uses (or can use) to justify that in the past they understood or meant themselves to be counting. Kripke (23) writes:

Am I supposed to justify my present belief that I meant [counting], not [quounting] . . . in terms of a hypothesis about my past dispositions?

I read his point as this: the *mere* description of one's dispositions to explain that one meant to count fits badly with the phenomenological impression we have of how people understand themselves to grasp concepts such as counting. One thing that indicates this is how odd the following speech would be:

I know how to count because I'm always disposed to answer the question "what is the number of items in that collection?" by pointing at each item in the collection one-by-one and uttering a numeral, starting with "1" when I point at the first item, and uttering the numeral that is the successor of the previous numeral I uttered each time I point to a new item, and to give the last numeral uttered when I point at the last item in the collection as the number of that collection.

If someone said this, she would raise the suspicion that she actually *doesn't* understand counting, but has instead memorized a procedure that she doesn't grasp at all. Or perhaps that she is describing a compulsion she feels rather than a course of action she takes under the guidance of a concept that she understands.²⁰ And we would suspect this precisely because her little speech doesn't indicate anything

about why she thinks a description of her dispositions is even *relevant* to whether she understands how to count.

I mentioned earlier that we often prove to someone that we know how to do something by actually exhibiting the appropriate behavior—e.g., by counting a collection: “You think I don’t know how to count? Well, watch this!” But in such a case the person isn’t offering her disposition to do such and such as a *criterion* of what counting is. Rather, her providing such illustrations to another person indicates that she presumes to already know that the other person knows what counting is, and she’s proving to such a knowledgeable person both that she knows *this*, and that she also knows what counting is. Both are achieved by an illustration of pertinent skills that she knows (and that she indicates that she knows) the other person will recognize to be pertinent skills.

4. THE FAILURE OF A STRAIGHT SOCIOLOGICAL SOLUTION TO THE RULE-FOLLOWING PARADOX

The answer that someone is disposed to give, when faced with a counting task, is a nonactual result that she would have exhibited if she had executed that task. Dispositions, so described, are severely limited in their ranges: they fail to satisfy Kripke’s first constraint because they can’t provide the needed psychological/individual grounding facts for every (possible) counting task, where such grounding facts are understood as sufficient for determining correspondence facts such as this: *Every collection of objects has one and only one cardinal number*. This is what makes it tempting to replace someone’s actual dispositions with dispositions that are counterfactually supplemented in various ways. Doing so is a very intuitively-natural response to the problem of limited powers because—as noted in the last section—we often excuse ourselves by saying that if we hadn’t gotten tired, we wouldn’t have made certain mistakes, or that we would have (successfully) undertaken a counting task that we have otherwise had to decline. In this way, we abstract the characterization of our knowledge of counting, and our ability to undertake it, from other factors that limit our execution of counting tasks in much the same way that what a Turing machine does is abstracted away from limitations of materials, tape, and energy. This is what gives intuitive force to the idea that Kripke’s

infinite constraint can be met by a counterfactual supplementation of our actual dispositions.

But such a strategy won’t satisfy Kripke’s justification constraint. And, indeed, that it’s even cogent to worry about whether the supplemented dispositions can do the job needed shows that the dispositions—*whatever they are and however they’re supplemented*—aren’t the standards by which correct counting is to be judged. If they were, there could be no question: whatever answer the dispositions yielded would be understood to be correct. But, regardless of what one’s dispositions to count look like, it seems that anyone can sensibly pose the question of whether his dispositions give the right answer; more strongly, it seems that it’s sensible to pose the question of whether one’s dispositions give the right answer. It’s bizarre for someone to say, when wondering if he has given the right answer after a count: “Well, it’s the right answer because it’s the answer I was disposed to give.”²¹

This invites looking elsewhere for the standards by which correct counting is to be judged. Furthermore, the contemporary scientific setting suggests that a focus on the dispositions of *individuals* to count is, in any case, largely irrelevant because the ability of individuals to count is one that’s very nearly never relied on. Most of the counting results we have aren’t based on any individual’s capacities, but instead on the use of scientific theories and instruments. Call the instruments involved “calculators.” Given the widespread use of calculators, it almost seems like the only relevant disposition that (most) individuals need is an inclination to accept results due to others, or due to calculators; all that’s required, that is, seems to be a disposition to *defer*. Indeed, the historical use of counting tools such as abacuses, counting boards, or knots on strings, indicates that something like this has been the case for many centuries. Our capacity to determine the numbers of many collections—and to estimate such numbers—doesn’t turn on anything reasonably described as the dispositions of individuals to count (unless, as just noted, we regard it as part of someone’s disposition to count that he will accept the results of instruments—calculators—that other people have invented and developed).

I stressed in section 2, during the discussion of Kripke’s justification requirement, that individuals feel justified in their choice of methods that they use to count. This is even true of the (confident) child who

will simply be puzzled by the adult's claim that everything she does is compatible with her knowing how to quount, rather than with her knowing how to count. Kripke uses this phenomenological fact against the dispositional account with great success because any subject's belief in his being so justified to do what he does isn't satisfied by the mere recognition on his part that he is disposed to undertake a process with such and such characteristics. This phenomenological fact, coupled with additional phenomenological facts, may seem to imply that the standards against which correct counting behavior is to be judged are social ones.

Consider, first, that the subject's confidence in her knowledge of counting is frail—as general discussions of scepticism make clear. Even if one is very confident of one's abilities, the right external circumstances can make one quite unsure of what one thinks one knows.²² This is not only because it's a fundamental part of one's understanding of any such task that one is "capable" of mistakes, and that these mistakes are seen as ones that are correctable by others or by instruments. It's also because it's a fundamental part of one's understanding of any such task that the magnitudes of the mistakes (that one is capable of) have, as it were, no lower limit. One therefore can be brought to become unsure of *anything* one thinks one knows. If how one were disposed to act or respond were *itself* the standard of one's behavior, such insecurity would make no sense. And *this* apparently shows that the standards in question are external to oneself, and located in the people around one—in one's community.

Indeed, it can be suggested that the normative elements that Kripke mentions as being part of the phenomenology, and that bear so negatively on the dispositionalist response, namely that (as cited in endnote 19), "Normally, when we consider a mathematical rule such as addition, we think of ourselves as guided in our application of it to each new instance," and "I follow directions," point pretty unequivocally towards the fact that a large part of our impression of being guided in the case of counting is that we feel we have grasped a rule *that others have taught us*. That is, that it's part of the *phenomenology* that the source of our confidence in what we do is that we take ourselves to have indeed learned it correctly (from others). And this explains why experiments can be so easily designed that undercut this confidence.

Everyone else isn't agreeing with me; and one possibility (therefore) is that I've failed to learn what I thought I learned—or that I've made a mistake—or worse, that I'm now experiencing the psychological effects of a medical condition.

The location of standards in the (dispositions of the) community one belongs to will also explain the result of a thought experiment, say, where someone suddenly becomes conscious of the subpersonal source of her counting ability. Suppose she suddenly starts to see bright red visualizations of numerals that are associated one-by-one with each object in a collection she's counting, and that the last one to so be associated with an item is the number she compulsively feels compelled to utter as the number of the collection. The natural reaction to this change in one's (conscious) psychology is the fear that perhaps the answers so forced *aren't* justified. This fear arises precisely because of the apparent detachment of the mechanism in question from the way we think of ourselves as having learned how to count. We learned it from others; but how can we be sure that this psychological compulsion we've now discovered to be operative in us is in accord with what we're supposed to have learned?

As I've been indicating, the foregoing considerations invite a straight *sociological* solution to the rule-following paradox.²³ This is the view that there is a pattern of grounding facts that indicates that the child is counting, and not quounting, but that the pattern is to be found (collectively) in the dispositions of the entire society within which that child has learned to count. Kripke (111) raises this option in passing, and notes that such a theory "would be open to at least some of the same criticisms as the original [individual-disposition theory.]"

Indeed, the straight sociological solution fails to satisfy Kripke's three constraints in *exactly* the same ways that the individual-disposition theory failed to do so. *The infinitude constraint*: The collective disposition of an entire community—howsoever that is defined in terms of collections of individual dispositions and collectively-available instruments—is still limited in its range, and still capable of performance deviations from correct counting. This is true even of *our* contemporary community—with its powerful computing instruments. Counterfactual supplementation of the actual dispositions of the community faces the same obstacles it faced in the individual case. *The*

justification constraint: A description of the dispositions of the community at a time doesn't indicate why that community is justified in computing new cases of counting the way it does—that by virtue of its (collective) dispositions the counting-concept it exhibits in the new case should be taken to be the same one that it previously exhibited in earlier cases. We do *not* think—even when we conceptualize counting tasks as ones we undertake collectively—that the way we do it is (by definition) the way it should have been done, given the way the community-practices were understood to operate previously. Finally, *the mistake constraint*: a collective solution offers no logical space for the required possibility of there being mistakes in calculation, not just by individuals, but even by an entire community. Attempts to meet Kripke's mistake requirement for collective dispositions are open to the question-begging concerns that were discussed in section 3.

5. ASSERTABILITY CONDITIONS, NOT TRUTH CONDITIONS

Let's turn, therefore, to Kripke's Wittgensteinian sceptical solution. A sceptical solution (Kripke, 66) concedes "that the sceptic's negative assertions are unanswerable," that is, it concedes that there is no pattern of grounding facts that can be used to answer the question of whether a subject is counting or quounting. Wittgenstein's solution (so Kripke argues) is to desert "truth conditions" for statements about a subject's understanding of numeration statements, and indeed "truth conditions" altogether. Because there is no pattern of grounding facts for someone's understanding of true statements (about numeration) that can meet Kripke's three requirements, there is nothing that determines necessary and sufficient truth conditions for statements about a subject's understanding those (numeration) statements. Instead, one substitutes "assertability conditions." "Assertability conditions" are the giving of conditions for when one is *licensed* or *entitled*—in a community—to assert (or deny) of oneself, or of others, that they understand counting. It's built into the notion of assertability conditions that a person belongs to a community in which such conditions are (recognized to be) in place.²⁴

Kripke suggests that the assertability conditions for A's entitlement to the claim that he means addition by 'plus' is that (roughly) he is

confident that he can give correct answers to new cases, subject to correction by others. And that he is also entitled (Kripke, 90), "again provisionally and subject to correction by others," to judge a new response by someone else as correct if that response is the same response that he's inclined to give. A denial that another understands counting is entitled if that person's answers deviate from one's own, and a denial of one's own understanding is entitled if one experiences insecurity about what the answers are or about how one is to go about finding them. Kripke (96) stresses that such assertability conditions have a point in our society precisely because (most) individuals in our society have similar dispositions.

As indicated above, "the community" is relevant in two ways to these suggested assertability conditions. First, a reference to the community occurs in the statement of the assertability conditions themselves: We are taken, that is, to recognize the provisional nature of our dispositions—that it's legitimate for *others* to correct us. Second, the community is relevant because what's on offer are assertability conditions; and assertability conditions themselves only make sense in the context of a *community* because they rely crucially on the notion of entitlement or license in a community to assert something or to deny it. Such concerns with entitlement or license lapse as relevant for isolated persons outside a community—what I'll call *Robinson Crusoes*.

According to Kripke's (86-88) interpretation of Wittgenstein, the incoherence of a lonely rule-following Robinson Crusoe is a corollary of the fact that such an isolated figure cannot be operating in a context with assertability conditions. When he takes himself to be counting, he acts "unhesitatingly but *blindly*" (Kripke, 87, italics his). The sceptical paradox tells us that if we restrict our attention to Robinson Crusoe alone, there are (Kripke, 89) "no truth conditions or facts in virtue of which it can be the case that he accords with his past intentions or not." Thus, there is no difference between Crusoe thinking he is following a rule and his actually following that rule; and therefore the idea that one can correctly or incorrectly follow a rule that one previously intended to follow collapses. As Wittgenstein (§202), writes:

To think one is obeying a rule is not to obey a rule. Hence it is not possible to obey a rule 'privately'; otherwise thinking one was obeying a rule would be the same thing as obeying

it.

Kripke (110) stresses that it *doesn't* follow that a Robinson Crusoe can't be described as following rules (for counting) correctly or incorrectly; it's that for someone to so describe Crusoe requires that person to treat Crusoe as belonging to his (the describer's) community.

Notice how the failure to make sense of normativity (that is, to provide a ground for Kripke's second, and especially third, requirements on any solution to the rule-following paradox) when restricted to the dispositions of the individual subject, is what's brought directly against the suggestion that a Robinson Crusoe can—independently of a community—be said to be counting. And notice also that Kripke's suggestion, that in our judging that a Robinson Crusoe is (or isn't) counting we must incorporate him into our community (and apply our assertability conditions, with our accompanying standards, to his case), follows directly from the fact that only in this way can the required normative elements in Kripke's Wittgensteinian assertability-conditions response to the rule-following paradox be enabled to operate.

6. A COMMUNITY OF IDIOLECTS

I intend to show—contrary to the foregoing—that a private model of rule following is possible. I undertake that task in stages, by presenting a series of differing Robinson Crusoes, starting in section 7, and continuing in the sections that follow. To prepare for this, I first briefly present in this section a thought experiment of a community that, contrary to the practices of *our* community, *does* fix the meaning of their words in terms of their dispositions. My intention is to explore some of the elements that can cause such an alternative practice to break down, and to illustrate some ways in which such a practice can succeed.

So imagine a community in which each individual really does intend each word that he uses, “table,” “count,” “add,” “owl,” and so on, to mean in each case exactly what he is disposed to take that word to mean. That is, by “dog,” a speaker means precisely those items he is disposed to characterize as “dogs,” by “table” those items that he is disposed to apply the word “table” to, and so on. It may seem that these suggested formulations trade on an inherent circularity that's clearly fatal to the proposal. After all, as noted, the meaning of “dog” is given

by a formulation involving the very word “dog.” It's true: there *would* be a circularity if the people in this community understood these formulations as providing them with *instructions* for using these words. But that *isn't* their understanding. Their dispositions—even if learned—are, and are recognized to be, automatic procedures that *dictate* answers to the persons with those dispositions; and so the formulation for the meaning of “dog” independently characterizes its meaning by treating it as applying to those items precisely when a disposition to so apply “dog” is operative. Call such languages “disposition-meaning languages,” and describe their words as having “disposition meanings.”

Notice, first, that for disposition-meaning languages, Kripke's justification requirement is met, and Kripke's mistake and infinitude requirements lapse. The justification requirement is met precisely because the invocation of someone's own dispositions *really does* justify that his current behavior is in accord with what he meant: people in this community always only intend to exercise their own dispositions, whatever they are. The mistake requirement lapses because there is no place in this language practice for a performance/competence distinction: a person's application of a word is always correct, no matter what his dispositions (and behavior) are like. One cannot perform “badly.” Lastly, no one in this community intends, in any case, to “add” or to “multiply,” etc., as *we* understand these notions: each only intends to perform “addition” and “multiplication” as his activities will be executed given his own dispositions. So only if someone's dispositions allow him to mean a notion of “counting” that applies to arbitrarily-large collections does he so mean it to apply.

One possible drawback of this disposition-meaning “language” is that, unless the people in this community have *exactly* the same dispositions, they will fail to have words in common; if there are deviations in their dispositions—however small—with respect to many words, the result will be a family of individual idiolects rather than anything that can be described as a “public language.” Furthermore, (i) they will recognize this fact, and (ii) they will recognize that there can be no translations from one idiolect to another—at least as far as a very large class of words is concerned. Under such circumstances, their cacophony of idiolects might seem to be useless for communication.

Such a pattern of idiolects won't be useless for (most) communi-

cation, however, if they deviate from one another in only small ways. For in that case communication failures will only occur sporadically. In general, therefore, how successful members of the population will be at communicating with one another will turn on how *much* deviation there is among their idiolects. Apart from this, and even if almost all their words involve substantial deviations, communication can succeed under certain circumstances, nevertheless, because translations can be possible.

Here's one way that could happen: Suppose B—a member of this community—has the word “disposition”; and suppose that her word applies *exactly* to C's dispositions and to her own. In such circumstances, B can recognize that how she means to understand her own words, nouns such as “table” and “orange,” verbs such as “run” and “multiply,” and so on, is via her own dispositions to apply such words. B similarly recognizes that this is true of C. B further recognizes that she cannot translate C's idiolect into her own directly: “table,” as C uses it, doesn't have the same application conditions as B's word “table” because their dispositions differ. Nevertheless, B can still translate C's word “table” like so: “table,” as C uses it, means in B's own idiolect, “‘table’ as C is disposed to apply ‘table.’” If everyone's word, “disposition” so applies to others' dispositions—if (this is a corollary of this assumption) they share the word “disposition” and if that word's application conditions exactly fit everyone's dispositions—then mutual translation of everyone's words in this community is possible.

Of course, we can imagine that *all* the words of all the speakers in this community deviate, including “disposition.” And, further, it can be that the way that B uses “disposition” causes the phrase “‘table’ as C is disposed to apply ‘table’” not to have the same application-conditions as C's “table” (because B's use of “disposition” doesn't correspond to C's dispositions). This, if recognized, would enable B to see that C and herself speak different idiolects that aren't mutually translatable to one another. Whether, of course, despite the impossibility of translation they could still communicate with one another would turn—as noted earlier—on how different the application conditions of their respective words were.

As the foregoing indicates, the impossibility of translating these idiolects to one another, and—more generally—the failure of such lan-

guages, at least for purposes of communication, doesn't inexorably follow from a wedding of the meanings of the speakers' words in a community with their individual dispositions to apply those words. This only follows given certain empirical facts about the idiolects of the members of that population.²⁵

We may nevertheless worry about the coherence of this thought experiment even in empirically fortuitous circumstances. Even in the best case, that is—when such individuals have uniform dispositions—can such individuals *consistently* apply words to objects? Why couldn't their dispositions be such that at one moment an object is labeled a table, and seconds later it's not to be labeled as such? As before, whether or not such idiolects can successfully enable communication turns solely on empirical facts about the dispositions possessed by each member of the community. Imagine that the uniformly-had dispositions to apply words in a community are extremely cautious ones. Perhaps it's a common disposition of individuals in this community not to apply the word “table” to an object until that object has been examined in such and such careful ways. Such dispositions would—to some extent—stabilize the application of words, and they would even stabilize (to some extent) the applications of calculational words such as “addition,” “multiplication,” and so on. Suppose, for example, that the members of such a community would only be disposed to apply such a calculation-word to their own computations if such computations had been undertaken and double-checked in such and such specific ways. Provided further that everyone had the disposition to give precisely the same answers in such cases (from our point of view: provided everyone was disposed to make *exactly* the same mistakes so that they *always* got exactly the same answers), then the idiolects would agree, and the applications of the words by the members of this community would be stable.

A provisional conclusion is this. At least as far as the discussion has been taken in this section, whether a disposition-meaning language—a language practice, therefore, for which Kripke's three requirements either lapse or are trivially met—is *practically implementable* turns purely on empirical considerations: exactly what dispositions the individuals in that community possess, and (importantly) in what ways those dispositions involve uniformities across the population. Of course, words also have applications—and some of them have applications to things

in the world, and to collections of such things. I haven't yet explored whether the success of such a language-practice also requires things in the world to possess certain uniformities that correspond in some way to the uniformities in the dispositions of the members of the population that are required. I also haven't yet responded to the challenge, that some philosophers might mount against the foregoing cases, that—regardless—these aren't cases where the individuals in question have *concepts* that their words can be taken to correspond to, and that they can be taken as applying to objects in the world. Rather, they just have tendencies to involuntarily express *noises* in ways that are more or less coordinated. The first issue will be dealt with in the course of presenting and discussing the various Robinson Crusoe cases, as I do in the following sections. The second issue will be addressed in the very next section.

7. ROBINSON CRUSOE IN FAVORABLE CIRCUMSTANCES

Imagine that Robinson Crusoe lives primarily on coconuts. Hunting for coconuts, however, is arduous (because he has to avoid various large predators). To help himself survive, he has cleverly taught himself to count the number of coconuts he has at a time, how many he eats in a day, and therefore how many days it is before he has to engage in a risky search for more coconuts.²⁶ He has also studied the territory he forages in, and he knows (for example) which trees have more coconuts and which have less—and therefore which trees are better for him to climb. In short, he's in a situation in which his counting skills are vital to his survival.

Included in the description of this Crusoe scenario, of course, is the apparently undeniable fact that collections of objects—e.g., the number of coconuts in that tree—have specific cardinalities. It's unsurprising, therefore, that in counting coconuts (and in counting other things, such as the number of days he can avoid hunting for more coconuts) it's important that Crusoe get the numbers he counts *right*. That is, it's important for Crusoe that when he counts the number of coconuts in a tree, that the number he gets corresponds to the number of coconuts actually in that tree.

Crusoe's applications of his other words, "coconut," "tiger," and so

on, must be judged on their adequacy in a similar way. Coconuts are food for Crusoe; he, on the other hand, is food for tigers. His word "coconut" is meant to pick out a particular set of items that are food for him; his word "tiger" is meant to pick out a particular set of items that it's best for him to avoid.

This use of "right" is different from the use of it that occurred earlier in this paper. What's apparently crucial to *this* use of "right" isn't whether Crusoe's counting practices are in accord with his earlier intentions, or with what he earlier meant by "count." What's apparently crucial to this use of right, to his answer being "right," is that his answers correspond to what's in the tree, so that his actions that result from his count (climbing one tree and not another to shake down its coconuts) yields the number of coconuts that he needs.

We can accept, I think, that Crusoe genuinely *understands* the words he has invented: "1," "2," etc., "coconut," "tiger," and so on. This is because he's given disposition meanings to these words: they are to mean exactly as his dispositions (at any time) incline him to apply them. However, given the notion of "right" just utilized, we can ask why so giving words disposition meanings doesn't lead to a complete disaster. Imagine, however, that Crusoe is in the following rather idyllic situation: (i) his dispositions to apply words are stable, and (ii) the way that things appear to him is always the way that they are. That is, if he thinks he sees a coconut, that's because he *does* see a coconut. I've put these conditions in what some might regard as an illegitimately tendentious way. What, after all, is a coconut *to Crusoe*, and relatedly, whatever does it mean to say that his dispositions to apply words are stable?²⁷

So let me put the matter differently. *We* live in a world where, as we like to put it, sometimes shadows look like coconuts, and sometimes coconuts look like cats. Furthermore, sometimes it takes a great deal of investigation to determine whether two things that we regard as alike in a particular way are in fact alike.

Imagine that Crusoe doesn't live in this kind of world. *He* lives in a world of *clones*. In our world, anything (pretty much) differs from anything else in its properties—apart from, I mean, their differing locations in space and time. Coconuts in our world vary in color, in shape, in size, in taste, and in numerous other ways. And this, of course, is

equally true of nearly everything else in our world. But in Crusoe's world, all coconuts are the same in all of their properties, all tigers are the same in all of their properties, all trees are the same in all of their properties, and so on. Everything in Crusoe's world falls neatly into one or another precise *natural kind*. Furthermore, Crusoe's senses are better than ours. Two things that are identical in their properties always look identical to Crusoe; two things that differ in at least one of their properties always look different to Crusoe. Finally, he can identify and distinguish objects at first glance; Crusoe *can't have* the experience that something looks one way to him at one time and a different way at a different time. It isn't possible for him to discover later that things that looked alike to him (or different) now look different to him (or alike). Finally, he never makes what we would describe as computational mistakes. If he *thinks* there are six coconuts in a tree as a result of counting them, that's because there *are* six coconuts in that tree.

In such a remarkable world, and with such remarkable epistemic powers, Crusoe *can* endow his words with disposition meanings, and enjoy a great deal of worldly success as a result. He *can* understand "7" to mean that cardinal number of a collection of items that he's disposed to apply "7" to. He *can* mean "coconut" to apply to exactly those things (all identical in their properties) that he's inclined to apply the word "coconut" to. Because Crusoe's idiolect is a disposition-meaning language, Kripke's three requirements either lapse or are trivially satisfied. I hasten to add that, because the infinitude requirement lapses, some would not regard Crusoe as *counting*, as *we* understand counting. Although I'm allowing that Crusoe never makes computational mistakes, I'm not assuming he can count arbitrarily large collections. Because there is an upper limit on the size of collections he is disposed to count, he has only a finite number of number words.²⁸

Nevertheless—or so I argue—Crusoe has a cogent language practice (a cogent idiolect) in which his words mean exactly what he is disposed to have them mean, and because of extraordinarily nice empirical circumstances, his idiolect is useful—indeed indispensable—to him. Furthermore, this language is useful for him despite there being no coherent application of the idea that he can apply his words wrongly.

Some would say that because there is no coherent application of

the idea of a wrong application for his words that he *hasn't* got a coherent language-practice. (And anyone who would say this would say the same of the community of idiolectal speakers that I considered in section 6.) They would deny, in particular, that Crusoe's words correspond to *concepts* that he can be taken to have: cardinality concepts, or the notion of a coconut, etc. Crusoe, as I've imagined him (so this opponent claims), doesn't have (real) concepts because his application of his concepts doesn't involve his being *guided* by such. Instead—and this is recognized by him—he has involuntary verbal *responses* to things in the world.

Three important points, however, support the claim that his concepts aren't quite the same as ours (in certain respects) rather than that he doesn't have concepts at all. First, there is mental content codified in Crusoe's use of his words. Although his understanding of his words explicitly gives those words disposition meanings, that doesn't stop his words from corresponding to collections of things in the world, and to collections that he experiences as similar, or as having certain cardinal properties, and not others. Indeed, this mental content in some cases *constitutes* his dispositions (e.g., he is disposed to treat such-and-such items as alike because of how he *experiences* them). Second, some of *our* concepts are just like his in similarly having disposition meanings. Consider, for example, the ordinary concept of *pain*. Notoriously, this is a concept to which "being wrong" doesn't seem to apply (at least in the first-person).²⁹ Lastly, and perhaps most importantly, nothing prevents Crusoe from using *his* concepts in exactly the ways we use *ours*: to describe things in his world, to reason about them, and to decide on a course of action on the basis of that reasoning. Just like us, he can decide to climb one tree rather than another (because of, say, the number of coconuts he takes each tree to have), or he can decide to challenge a group of dangerous predators because there are only two of them, and not three. These striking facts about Crusoe's abilities correspond to similar facts about our concepts and our abilities: those concepts of ours to which "correct" or "incorrect" don't seem to apply, such as "pain," seem nevertheless to be concepts that can be used together with our other concepts (to which "correct" or "incorrect" do apply) to describe aspects of our world, to reason, and to make decisions about courses of action.

One last observation: It would be a mistake to imagine that Crusoe—in such epistemically favorable circumstances—would see himself as omniscient. Certainly he would lack the notion of a “mistake” in one sense of that word: He wouldn’t understand how it’s possible to think something is a coconut, and yet for it not to be such. But despite this, he would still understand what it means to be ignorant. He would have experienced (indeed, he would be quite familiar with the experience of) not knowing how many coconuts there are in a tree because he hasn’t counted them yet, or not knowing whether there is a tiger around the next bend because he hasn’t gone round that particular bend yet.

There is a subtlety here that I should at least indicate by distinguishing between two kinds of cases. One kind is this: Crusoe might face a collection of items—like a pile of sand—that he has no disposition to attempt to count. In this case, he is neither disposed to apply any of a large number of counting words or to deny applying them: as far as his idiolect is concerned, there is no answer to the question how many grains of sand there are. He’s not ignorant of the answer; there isn’t one. But there is a different kind of case, where he feels he would be able to either apply a word or not apply it if he chose to—for example, he may be disposed to deny that “1,” “2,” or “3” apply to the collection of coconuts in a tree. As a result, he knows that there are more than three coconuts in that tree; but without counting (without engaging his dispositions to supply a cardinal number to the collection of coconuts in that tree), he doesn’t know exactly how many coconuts are there.³⁰

In the fortunate circumstances our first Robinson Crusoe finds himself—fortunate both because of the idyllically simple world he lives in, and because his dispositions perfectly match what there is to be found in that world—he has no need of the concept of the “correct application” of a word. In this respect his resulting idiolect is different from the language we take ourselves to speak. It is, however, recognizably a language, and recognizably a useful one. The Crusoe 1 thought experiment has therefore already tamed the scope of Kripke’s Wittgensteinian argument against the private model of rule following. Leaving aside the slight terminological abuse of describing this as a case of private rule-following, the unexpected result is this: In the right empirical

circumstances, there are useful private languages—not because in such cases there is a difference between the isolated individual following a rule and his only thinking he does, but—because his dispositions and his world are such that there being no such difference doesn’t matter.

8. A ROBINSON CRUSOE WITH TWO SETS OF DISPOSITIONS THAT HE HAS CONSCIOUS ACCESS TO

Let’s now imagine a second more realistic Crusoe 2 who is in a world that’s the same as the one our previous Crusoe 1 was in: picturesquely put, it’s a world of different kinds of clones. As before, there are food-stuffs (“coconuts,” “shrimp,” etc.) and predators (“tigers,” “boars,” etc.); but every shrimp is identical in its properties to every other shrimp, and so too for boars, trees, blades of grass, pebbles, etc.

Crusoe 2’s dispositions, however, have certain instabilities. Things sometimes seem one way to him, and sometimes they seem another way. Sometimes, when the light is different, or Crusoe 2 is tired, or if he has made a quick decision, or when he has had too much of his home-brewed liquor to drink, what he is disposed to call a “coconut” is something he will be disposed to later call, under other circumstances, “a squirrel.” Sometimes, what he is disposed to describe as a “tiger” is what he is disposed (after staring a little longer, or after climbing a tree and waiting to see what the thing will do) to later call a “shadow.”

A first (and rough) way of putting the difference between Crusoe 2 and Crusoe 1 is that it would be valuable for Crusoe 2 to learn to recognize when he can trust his dispositions and when he can’t.³¹ There are two possible cases. The first is the less realistic one where Crusoe 2 can distinguish two internal states that it’s possible for him to be in. The first state involves one set of dispositions, and the other a differing set of dispositions. For the sake of simplicity, I’ll call one state “being rested,” and the other “being tired.” The second, more realistic case, is one in which his dispositions can vary over time, but he’s unable to classify those dispositions in any neat way, or even recognize (much of the time) which state he’s in and whether his dispositions have shifted as a result. I take up the second kind of case in section 12.

In the first case, Crusoe 2 has a choice between two possible disposition-meaning languages. Recognizing, that is, that the two

states cause him to give his words different disposition meanings, he can be brought to recognize that one set of dispositions (one disposition-meaning language) is superior to the other. He can decide that one is “trustworthy,” and that the other isn’t. Imagine, in fact, that his dispositions—when he is rested—match the world in the way that the Crusoe 1’s dispositions matched his world. Crusoe 2’s dispositions when he is tired, however, deviate. It seems that Crusoe 2 can notice this fairly quickly. When he’s tired, his word “coconut” has his tired disposition-meaning of “coconut.” But he soon discovers that tired-disposition coconuts are not all alike. Some are tasty; but others try to bite him. According to the rested-disposition meanings of his words, however, only some of these items are coconuts: the ones that try to bite him are squirrels.

Let’s describe Crusoe 2’s discovery in a slightly different way. He has discovered that he has a choice between two different disposition-meaning languages, each based on a different subset of his dispositions. And one is tempted to say—indeed, this is just how I put it before—that he chooses the more “trustworthy” set of dispositions to underwrite the meanings of his words. I also spoke of his rested dispositions “matching” the way *the world is*. But what does all this talk mean to Crusoe 2?

In asking this, I’m not asking the practical question of how Crusoe 2 manages to make a recognizably better choice between the two languages. It’s relatively obvious how he manages that because it’s clear how one set of dispositions keeps getting him in trouble, and the other set doesn’t. We might be tempted to put the point this way: he *thought* something was a coconut, but when he sat down to eat it, it attacked him, and that’s when he realized it wasn’t a coconut. But this isn’t right. To see this, let’s use ordinary quotes to describe words in the language *we* use to describe Crusoe’s situation; let’s use r-quotes for his rested language, and t-quotes for his tired language. We don’t want to say he thought the item was a “coconut,” but that he discovered it wasn’t. This would be to use *our* word, a word he hasn’t access to.³² But, he didn’t discover the item wasn’t a ^tcoconut^t, because it *is* a ^tcoconut^t. We *can* say he discovered it wasn’t a ^rcoconut^r. But it’s not that he *thought* it was a ^rcoconut^r. He wasn’t using his rested-disposition language when he uttered the sound, kō’k(schwa) nut’; he was uttering

his word ^tcoconut^t, which applied correctly to what he applied it to.

Furthermore, the trouble that Crusoe 2 has gotten into seems to be something that can be expressed in either language that’s available to him. According to how Crusoe 2 can speak in his tired-disposition language, the trouble is that some ^tcoconuts^t bite and some don’t. But this doesn’t imply that they’re not ^tcoconuts^t, or even that they’re ^tdifferent^t. After all, they are ^rdifferent^r, but that doesn’t make them ^tdifferent^t. How can these *differences* as described in the two languages be separated by Crusoe 2 so that one is pertinent to what the (biting) object *really is*, and the other isn’t?

We’re tempted to say: The rested-disposition language *really* describes coconuts as they are. After all, Crusoe 2’s word ^rcoconut^r agrees in its extension with our word “coconut,” and our word gets it right. But what’s wrong with the extension of ^tcoconut^t? *We* think some ^tcoconuts^t are coconuts and other ^tcoconuts^t are squirrels. But so what? The words in each of Crusoe 2’s meaning-disposition language have different scopes—they range over different collections. And there is nothing wrong with *this*.

We’re tempted to say in response: The set of collections of objects as categorized by Crusoe 2’s rested-disposition language accords (better) with the real resemblances among the objects so collected together; and the set of collections of objects as categorized by his tired-disposition language doesn’t. One is tempted to say more: The rested-disposition language *matches* its words to the natural kinds that are in the world, and the other doesn’t. Coconuts and squirrels simply don’t belong together in a kind.

There are two issues here. The first involves the question of what gives *us* the right to make this claim. This is an issue to be probed later in this paper. The second issue is that, in any case, Crusoe 2 hasn’t access to these considerations. He’s got *two* phrases ^treal resemblance^t and ^rreal resemblance^r. That is, each of his languages—by its own lights—characterizes the real resemblances in its own way. He also has two phrases that correspond to our words “natural kinds.” Each of his languages—by its own lights—characterizes the “natural kinds” differently. With what words is Crusoe 2 supposed to make sense of the idea that the words in one disposition-meaning language are “more accurate” than the ones in the other disposition-meaning language?

Notice the point. It isn't (or it isn't *yet*) that *we* can't make sense of the idea that one language better fits the world than the other does. As I've been indicating, we apparently do just that by comparing the extensions of his words to *ours*. The issue at the moment is this: how is *Crusoe 2* supposed to make this distinction—for himself—have content that goes beyond the already stated fact that one language enables him to navigate the world better than the other one does?

Let's table the question of *our* right to describe *Crusoe's* words as fitting (or not fitting) the world by imagining that, from *outside* both of *Crusoe 2's* possible meaning-disposition languages, God has a language each word of which, because of *His* powers, refers to a collection that contains only clones; imagine too that *His* phrase "natural kind" refers only to the collection of those collections of clones. (Don't ask how *God* makes *His* language do this; He's *God*, after all, and a supernatural ability like this one comes with the territory.) Using *His* own language, *God* can describe *Crusoe 2's* situation this way: There *are* the ways that things are, and there *are* the ways they really resemble each other—by belonging to the same natural kinds. The ways that things are, and the ways that they really resemble each other, *affect* *Crusoe 2*. One of *Crusoe 2's* languages—the rested meaning-disposition language—matches the world in these respects, the other doesn't; and so when *Crusoe 2* chooses the rested meaning-disposition language he can navigate what's in the world more successfully than he can if he chooses the tired meaning-disposition language. In the first case he categorizes and groups things by the ways that they are, and the ways that they resemble each other, and in the second case he fails to do this.

This way of putting things (*God's* way of putting things), as noted, isn't available to *Crusoe 2*. Despite this, *Crusoe 2's* choosing the rested meaning-disposition language over the tired meaning-disposition language is rational because he can recognize the superior choice by how it improves his ability to make his way in his world. He sees that he is enabled to succeed when adopting one language in ways that he isn't so enabled when he adopts the other language. He can't capture this difference, as *God* can, by comparing the two languages directly to the world, and seeing that one "fits" the items in the world better. For what can *he* mean by this use of "fit"?

Notice what *Crusoe 2's* problem is. It's that he has—let's say—the

ability to understand that one language is superior to another because it enables his success in the world. But *this* is not a necessary and sufficient condition for a language better "fitting" the world. He has, further, the ability to imagine that one language will be superior to all other possible languages (for all time), and in every way, for enabling him to succeed in the world. But this too isn't a necessary and sufficient condition for a language "fitting" the world. I'm suggesting that it isn't possible for *Crusoe 2* to understand what a language fitting the world can mean. (And, I'm also suggesting, that *Crusoe 2's* problem with understanding how *God* can make *His* language fit the world is deeper than it initially appears. Perhaps, *we* don't understand what this can mean either—this is a question I dedicate section 15 to probing further.)

Why can't *Crusoe 2* understand the idea of a language fitting the world? The reason is this. By assumption, his language is a disposition-meaning language; relatedly, his concepts are disposition-meaning concepts. That means that we are to describe his understanding of any concept in terms of his dispositions to apply that concept. But it's precisely this that goes missing in the case of his view that one language is superior to another by virtue of its better "fit" with the world. That's not how he responds to the comparison of two disposition-meaning languages. *He* responds only in terms of which language enables him to better succeed in the world, and that isn't the same thing as the words of such a language better fitting the world.

One thing *Crusoe 2* *can* do is adopt one language, and fault the disposition meanings of the other language in terms of it. We might think that the way he should indicate this is to say, using the rested meaning-disposition language: "I was tired, so I *thought* I was carrying a coconut, but I wasn't. It was a squirrel." In saying this, he would be speaking in a way—as we've already seen—that's entirely analogous to how he could speak if he had instead adopted the tired meaning-disposition language: "I was well-rested, so I *thought* I was carrying a squirrel, but I wasn't. It was a coconut."

It's important to realize that, for *Crusoe 2*, this apparent way of speaking about being mistaken about the *objects*—"I thought the object was a B, but I was *wrong*." "I made a *mistake* about what kind of object I was dealing with"—isn't *required*. It may, in fact, be pretty

unnatural for him. Because Crusoe 2 is aware that there are two meaning-disposition *languages* he can choose between, and that he has less success with one than with the other, he can instead describe the situation metalinguistically. He can say instead: I used the wrong word *coconut*. In doing so, he still describes himself as having made a mistake; but now it's a mistake about the word he used, and not a mistake about something in the *world*. Or, equivalently, he can say, "I used the wrong dispositions when I uttered the sound, *kō'k(schwa) nut'*." In the second case, he thinks of himself as mistakenly employing the wrong dispositions for the single sound *kō'k(schwa) nut'*. In the first case, he instead thinks in terms of there being two different words, individuated by different dispositions, and he has used the wrong word. For Crusoe 2, the difference seems to be merely terminological.

To repeat, Crusoe 2 can choose one language over the other only because it's superior in enabling him to navigate his world—and not because (as God can) he can see that one "fits" the world in a way that the other doesn't. His doing badly if he adopts one language, and his doing better if he adopts the other, is an objective fact. But it's not a fact *he* can express by a comparison of how languages purport to contour the world to how the world itself is contoured.

We've returned to the concept of *fit*. Some philosophers might think that Crusoe 2 can think of *fittingness* as an *explanation* for why one language is superior to the other: the superior language fits the world better—that's why it's better for navigating his world. Crusoe 2—so some of these philosophers might say—is engaging in an inference to the best explanation for his success in the world: his success must be due to the fact that his terms are picking out things in the world in the way that those things are. His success must be due to the fact that the ways that his terms carve the world matches the way that the world is carved—at its "joints," as it were. He doesn't *see* that this is so; he *infers* from his success that it must be so.

We still have the problem we had earlier: what can Crusoe 2 *mean* by these metaphors: "fits better," "carves the world at its joints" when he gives this kind of explanation? How does the meaning of these phrases go beyond the evident fact that he can better navigate the world by means of the language he has chosen? One thing these metaphors *can't* mean is this: Given the "correct" description of the

world, the superior language describes the world in the same way. It can't mean this because Crusoe 2's notion of a "correct description" of the world can only be a description that's a better *fit* than any other description (it's one that succeeds in carving the world at its joints). And now it's clear that the same problem has arisen again: what are these metaphors supposed to mean to him?³³

Suppose that Crusoe 2 quounts, rather than counts. As we've seen, *God can say*: Crusoe 2 will do badly in his world because he will systematically treat collections as identical in cardinality when they're not.³⁴ On the other hand, if there were someone else on the island with him (*Friday*, say), and Friday counted rather than quounted, Friday would soon discover that he could regularly cheat Crusoe 2 out of coconuts. In commonly used language, he could use Crusoe 2 as a *money pump*. Suppose that both Friday and Crusoe 2 have exactly the same rested, and tired, dispositions. If Friday (wisely) chooses the rested disposition-meaning language, and Crusoe 2 instead (foolishly) chooses the tired disposition-meaning language, then Friday can take advantage of Crusoe 2 in *all* sorts of ways. Friday can describe this from the vantage point of his own language by saying that because Crusoe 2 thinks he has 5 coconuts, he's willing to trade them for Friday's 5 sardines (but, according to Friday, Crusoe 2 is actually trading 57 coconuts for 5 sardines). Friday can *say* this; but *his* authority for his claim can only be that he can successfully take advantage of Crusoe 2. He hasn't God's reasons for his claim (even though his and God's languages are word-for-word translatable in the sense that Friday's application of his words matches God's application of corresponding numerical terms). God knows what He knows, by assumption, because He makes His words fit the world metaphysically—or, perhaps because He made the world to fit His words. Friday hasn't got God's supernatural powers: so Friday can only know that he's doing far better than Crusoe 2 is.³⁵

Let's return to the isolated Crusoe 2 choosing between his two possible meaning-disposition languages. To repeat: the evident fact that Crusoe 2 does better with one language than with another is what enables him to make a choice. And, having chosen, he can now express his superior choice *in* the language he has chosen by, as we saw earlier, either faulting his dispositions that his use of words shouldn't rely on, or by faulting the language (with particular disposition mean-

ings) that he has inadvertently slipped into. Doing so, however, doesn't (and can't) rely on Crusoe 2 knowing the metaphysical fact that terms in the superior language carve reality in a superior way—e.g., more accurately—than terms in his earlier language did.

Nevertheless, the foregoing suggests that in the right circumstances, Crusoe 2 can realize that it's appropriate for him to trust *some* of his dispositions, and not other ones. Necessary to this ability, of course, is that he can identify (some of) his internal states. *Given* that ability, his favoring some of his dispositions over other ones *isn't* an arbitrary matter that depends on his will alone (*pace* Kripke's Wittgenstein). Rather, it turns on objective facts about the success value of the different disposition-meaning languages that he can choose between. I explore this last point further in the next section.

9. SUCCESS AND FAILURE

I've suggested that Crusoe 2 can distinguish his dispositions on the basis of the subsequent success he experiences when relying on them. Crusoe 2, so I'm assuming, can distinguish and grade his well-being subsequent to events that affect it. And, I'm also assuming, he can recognize how his well-being is affected by those events. It's not necessary to assume that Crusoe 2 is entirely objective about this, or that he will always be right. It is necessary, however, to assume that he's pretty good at recognizing when what he has eaten has made him ill or when an animal is dangerous, or has hurt him, and so on. In allowing that he can do all of this, I am, of course, also assuming that Crusoe 2 has a number of dispositions that enable him to learn from his experiences.

One issue that these assumptions raise—right at the start—is that different Crusoes can have different “systems of values.” The Crusoes I've imagined in previous cases (and will imagine in the cases to come) are straightforward creatures with values that we find easy to understand: they want “creature comforts,” they want to survive, they want to avoid pain and discomfort, they want good meals (subject to the constraints of their island resources). But other Crusoes are possible: ones who disdain creature comforts, who wish (masochistically) to be eaten by tigers, who embrace pain and discomfort and death. The appropriate notion of “success” seems relative to the particular values of

the Crusoe in question.³⁶ Consider again the case where Crusoe has a choice between two languages: the tired disposition-meaning language and the rested disposition-meaning language. Surely it's the case that subsequent success for a Crusoe with certain (self-disaffirming) values would turn on his choosing the tired disposition-meaning language rather than the rested one. For in that case, he would facilitate his own death—which, let's say, he desires to occur in a certain way (e.g., by an accident brought on because, at least in part, of his own incompetence).

There is no denying the relativity of a Crusoe's success to what that Crusoe wants. “Success” must be, therefore, at least partially relativized to the aims and desires of the Crusoe in question.³⁷ Two points, however, have to be made about this. The first is that the aim of this paper is to introduce a series of Crusoes with a broad range of dispositions, and in a family of environments, where such Crusoes can engage in cogent private rule-following. I intend, in the previous case, and in the cases to follow, to root that cogency in a trajectory of modifications of a Crusoe's own dispositions that lead to “successes” that the Crusoe can be aware of, or at least that he can respond to, and that—at the same time—are objective. I'm *not* claiming that we can attribute to a Crusoe any set of dispositions whatsoever, and expect that such a Crusoe will be able to engage in cogent rule-following. That clearly isn't true. So in this paper I'll be sticking to relatively simple cases: ones where a Crusoe wants to survive, wants creature comforts, etc., and I'm going to defer “unhealthy” Crusoes for later work.³⁸

The second point is this. Although “success” must be relativized to a Crusoe—more specifically, to his values, still more specifically, to what he wants—at the same time this success must be objective. It is, that is, an objective fact whether Crusoe really is better off as a result (relative to the standards provided by his values) when he takes there to be so-and-so many coconuts and not a different number, and so on. It's an objective fact whether he is better off grouping the creatures in his world in this way rather than in that way. These objective facts are relative to what it is that Crusoe wants; but they are objective nevertheless.

This answer seems, however, to raise two further issues. The first is that such objective facts, about whether Crusoe is better off or not,

seem to be ones that must be expressed in a language that really describes the world as it is. At the same time, it also seems that such objective facts must be rooted in the propositional attitudes of Crusoe: that *he* wants such-and-such, and that *he* doesn't want so-and-so.³⁹ But any such description of Crusoe's propositional attitudes must—so it seems—be couched in some language or other. As indicated in section 8, since Crusoe articulates to himself that he wants a coconut, and that he wants to avoid a tiger, it seems that such attitudes must be couched in whatever language he (Crusoe) is using—in particular, his words “coconut,” “tiger,” etc., seem needed. But his language doesn't seem to allow a description of the appropriate failures. After all consider Crusoe 2 using the tired disposition-meaning language: he *wanted* a ^tcoconut^t. And he got a ^tcoconut^t. Where's the failure? If we describe such attitudes, using an objective language (God's language, say), then we are no longer expressing Crusoe 2's propositional attitudes as he would express them, and perhaps not even as he would recognize them.

The first move to make is to deny that Crusoe 2's success is to be couched in terms of “coconuts” and “tigers.” Rather, it's to be couched in terms of the satisfaction of his more primitive desires—hunger, fear, safety, and so on. It's in terms of these concepts, and not in terms of words such as ^tcoconut^t or ^rcoconut^r that his successes and failures are to be indicated.

But this move seems to just push the problem back a step. There are two options. The first is that Crusoe 2 has words for these psychological states as well: ^thunger^t, ^rhunger^r, ^tfear^t, ^rfear^r, ^tcomfort^t, ^rcomfort^r, and so on. And so, as my r- and t-quotes indicate, it seems that the same issue arises for these words that arises for his coconut-words and his tiger-words. The second option is for Crusoe 2 not to have words for these psychological states. But it's unclear why this second option would enable an escape from the concern about the language of Crusoe 2's propositional attitudes. For in any case, he certainly has methods for recognizing that he is hungry or fearful. And when he's tired, he might think he is hungry when he is not, fearful when he is not, or in pain when he isn't. So it looks like we still have the same problem because even without explicit words, Crusoe 2 is still saddled with two sets of concepts.

This problem can be neatly solved if we attribute to Crusoe 2 a certain introspective access to his own states—so that he doesn't think he feels hungry when he doesn't feel hungry, fearful when he isn't, or in pain when he isn't. In doing so, Crusoe 2's psychological self-impressions are linked with *certain* of his psychological states. Crusoe 2, at least in these respects, is always correct about certain of his own psychological states; and that means that if his success and failure is couched in terms of *these* concepts, pain, discomfort, feeling hungry⁴⁰, and so on (whether he articulates them or not), the problem is solved. For Crusoe 2's propositional attitudes, his desires and hopes, when couched—not in terms of coconuts or tigers but in terms of pain, discomfort, and hunger, are objective: they are the same notions that could be used by God if He described Crusoe 2's propositional attitudes. Furthermore, these concepts (or the words that indicate them) are rigid: they don't shift as Crusoe shifts in his languages.⁴¹

This solution doesn't pose the danger of creating an incoherence in the Crusoe 2 thought experiment, or in its descendants; it also poses no danger of showing that Crusoe 2, counter to my claim, fails to engage in cogent private rule-following. Rather, the danger it poses is one of making Crusoe 2 (and the descendant Crusoe cases discussed in later sections) so artificial that the cases of private rule-following they allow are too *recherché* to take seriously. I address this concern in the next two paragraphs.

Let's start with one important qualification. For this solution to work, it isn't required that Crusoe 2 have *infallible* access to his introspective states: pain, fear, etc. Perhaps he does get them wrong once in a while, or every so often. (And we know that someone can be affected by certain diseases in such a way as to get them wrong systematically.) If Crusoe 2 doesn't have infallible access to his own psychological states, we must allow that he has two sets of concepts (or words), e.g., ^tfear^t and ^rfear^r. If these paired concepts deviate too far from one another (and from his internal psychological states in particular ways), a Crusoe 2's capacity to engage in private rule-following will collapse. But as long as these words *largely* agree in their extensions, correspond enough with his internal states, and thus largely agree with God's words, they will provide Crusoe 2 with the ability to recognize success, and consequently to settle on disposition-meaning

languages that enhance that success. And, in describing his wants, outsiders (God, ourselves) will be able—with only a little inaccuracy—to use their (His, our) own words “fear,” “hunger,” etc., to describe the relevant states.

Next, it’s worth noting that more or less successful introspective access to one’s own creature-comfort level seems to fit very well with our self-ascribed psychology. Even though contemporary pop psychological literature is saturated with fallout from studies that show how deluded we can be about our own motivations, and how poorly we often judge our capacities and our performances, it is still true that we seem—broadly speaking—to be aware of when we are in pain, and in discomfort. Indeed, we are pretty good, when it comes to basic creature comforts, at recognizing when we are better off (or worse off) than we were before. I stress again: our skills in this area don’t have to be all that good to enable successful private rule-following. More significantly, our varying abilities in this respect translate rather directly (leaving aside the factor of luck) into how well we can engage in private rule-following. Not only is that all that’s needed; it’s the best one should expect.

My claim is that the qualification of the paragraph two paragraphs before this one, and the observations about our ordinary views—just sketched—about our capacities to introspect (some of) our psychological states are enough to blunt the worry about the private rule-following cases like that of Crusoe 2 being too *recherché*.

10. CAN CRUSOE USE THE WAYS THINGS ARE IN THE WORLD AS A STANDARD FOR HIS WORDS?

Some philosophers might try to challenge what I’ve claimed in section 8 this way: Surely Crusoe 2 is comparing the two disposition-meaning languages—not by asking himself which set of dispositions enable him to better navigate his way around the world, but—by the very different question of which sets of words it is that (best) captures what’s out there. I’ve already suggested that this *isn’t* possible for Crusoe 2 but the view that somehow this *is* possible is so seductive (and popular) that the view deserves further discussion. I’ll approach it in a different way.

It’s certainly natural to think that Crusoe 2 recognizes that he wants his words to refer to collections of things in the world that *are* alike, and not merely to collections that he is *disposed* to treat as alike. But how is he to understand his words as doing this? More strongly, how is he to *make* his words do this? In the last section, I entertained the idea that God has a language where each word fits a natural kind perfectly, and that His phrase “natural kind,” as well as phrases like “the same kind as,” and so on, pick out, and operate with respect to, all and only the natural kinds that there are. We don’t need to explain how God’s language has this property; He’s omnipotent (and so we can stipulate this ability of Him without—and this is convenient—*understanding* how that’s possible). We *do*, however, need to explain how a Robinson Crusoe can do anything similar.

So we are now engaging in a thought experiment with Crusoe 3. His situation (and his dispositions) are identical to those of Crusoe 2 except that he intends his words not merely as his dispositions incline him to so apply those words, but as applying to certain items in the world and not to others.

To make this work, we need to explain how Crusoe 3 is to understand *his* words “coconut,” “tiger,” “natural kind,” and so on, so that *they* mean what he—according to this suggestion—intends them to mean. The first step, it will be said, is that Crusoe 3 can’t understand the standards of his words to be as his own dispositions (or a subset of them) would incline him; he must instead understand the standards of his words to be the way things are in the world. But how is *he* supposed to manage this? All he has, after all, to press his words correctly onto the world (as it were) are his actual dispositions to apply these words—that’s as far as his *natural* powers extend. How is the switch to be made from his *actual* dispositions doing this (and thus his dispositions being the *de facto* standard for what his words mean) to the things in the world being the standard instead? How, in other words, does the way the world is the standard for the correct application of his words manifest itself in his intention? He can’t, for example, think or say: “the standard for the application of my word ‘coconut’ are the *real* coconuts.” After all, “real,” as in “real coconuts” or “real world” is one of *his* words.⁴²

Crusoe 3—like Crusoe 2—must take himself to be *relying* on his

dispositions to group things as belonging to the “same kind,” or as “similar,” whenever he applies the word “coconut.” What other option is there? And this means that he applies the word (concept) “coconut” to whatever it is that strikes him as similar to *that*. Nevertheless, the suggestion is that he *understands* the meaning of his word “coconut” not as applying to whatever it is that *strikes him* as similar to coconuts, but as applying to whatever it is that *is* similar to coconuts. (But, given that this is to be his *understanding*, “is similar” must be *his* phrase as well.)

Notice that there are two problems we’re facing here. One is how Crusoe 3 is supposed to *make* his words obey the world’s standard. The second is how Crusoe 3 (using his words and his concepts) is even supposed to be able to *think* or *understand* his words as the suggestion is requiring him to be able to do.

Let’s leave these puzzles aside for a moment. Notice that one necessary implication of this way of thinking that’s being imputed to Crusoe 3 is that he will treat his specific characterizations of items as “coconuts” (based as they can only be on his dispositions to apply “coconut”) as *only defeasibly* applying. He will have to regard himself as capable of *mistakes*. This notion of “mistake” is importantly different from the one that Crusoe 2, having decided on the rested disposition-meaning language, had. Crusoe 2 describes his tired dispositions as mistaken by using his rested dispositions as the standard against which his tired dispositions are to be judged: he *thought* such-and-such was an owl (he means: “my tired dispositions indicate that this is an owl, but those dispositions don’t determine what ‘owl’ means”); but it wasn’t an owl (he means: “my rested dispositions reject this as an owl, and those dispositions do determine what ‘owl’ means”). Crusoe 2 recognizes mistakes as using the wrong words (or the wrong dispositions with the right words)—using words or dispositions that are less valuable than ones he could have used. Crusoe 3, however, is supposed to be intending something different: *the world* is to be the standard, not any subset of his dispositions. And this means that his mistake isn’t in his choice of language, but is one about the *objects themselves*. He got something wrong about the *objects*, and how those objects should be grouped.

So we must return to our original questions. First, how does Crusoe

3’s intentions enable his words “coconut,” “similar,” “natural kind,” and so on, to bend to the “world’s standard” (to, in other words, “match” the particular natural groupings that are in the world)? *Something* must be enabling this to happen, and so this way of posing the question is natural: What are the facts about the world (and Crusoe 3) that determine that his word (concept) “coconut” refers to this collection (*coconuts*), and not that collection? What are the facts about the world (and Crusoe 3) that determine that his word (or concept of) “similarity” groups things in collections that correspond to kinds as they are in the world, and not in some other way?

Second, how can Crusoe 3 even get his thinking (since it’s in the medium of his own words and concepts) to be *about* “things as they are,” independently of his own dispositions to use his words and concepts, so that he can even *entertain* the suggestion that the standards for his words are to be things “as they are,” as opposed to things as his dispositions determine him to think of them?

Crusoe 3’s supposed understanding of the meanings of his own words has created a space between what all his words “coconut,” “natural kind,” “real,” etc., are to “actually” refer to, and his resources for attaching them to items in the world. And the questions (I keep repeatedly pressing) focus on the supposed facts that are needed to determine—independently of Crusoe 3—what these words refer to. The problem is simply that Crusoe 3’s dispositions *exhaust* his resources in determining what his words refer to; they exhaust what he can even think of his words as referring to. How can anything else help? (How can anything else change this?)

What a strangely large number of philosophers in an influential tradition in metaphysics are tempted to say at this juncture is that *the world* does the job for Crusoe 3 because objects in the world *themselves* sort into natural kinds, and “coconut” as used by Crusoe 3 is “stipulated by convention” to refer, if it refers to anything, to a natural kind.⁴³ One way it might strike philosophers that this could work is this: As long as Crusoe 3 can be sure that a *couple* of items are similar (belong to the same kind)—and surely *that’s* not too much to demand⁴⁴—and so long as items in the world behave nicely enough (sort into simple enough kinds, as I’ve stipulated to be the case with respect to the items in the worlds that all the Crusoes so far discussed live in), then Crusoe 3’s

word “coconut,” for example, will refer (by stipulation and by the fact that he’s got a couple of items he can be sure are “the same”) to all and only the coconuts. The natural kinds in the world *themselves* operate as “reference magnets”: they *pull* the references of Crusoe 3’s words onto the natural kinds so that his word “coconut,” for example, refers to all and only the items that belong to the same kind as his pair of sample coconuts do.

This suggestion *so* misses the point. Regardless of how neatly items in the world are organized (and to repeat, worlds don’t get much neater than the ones I’ve described the Crusoes as living in), we still need to explain how *Crusoe 3* manages to get his word “coconut” to cover all and only the coconuts. It isn’t somehow built into his words—it isn’t built into *any* words—that (of course) they won’t refer to irregular portions of reality. (These are *his* words—how was such a condition built into his words, and by whom?) Collections of items in the world—no matter how naturally they belong together—don’t “semantically throb” or “semantically radiate” in some magical way so that Crusoe 3’s words can pick this up (like little radios) and so refer just to *those collections* instead of to irregular portions of them.⁴⁵ As noted, if it’s suggested, instead, that Crusoe 3 manages the trick by relying on his prior concept of “natural kind,” “regular portion of reality,” or some such other fundamental notion, the same questions can be raised about *those* words or concepts.⁴⁶

The foregoing explains why one can’t claim that it’s a “convention” of Crusoe 3’s language (or of any language, for that matter, apart from God’s) that his concept, “same kind,” applies to the natural kinds in the world; no more can one claim that it’s a convention or a stipulation that Crusoe 3’s notion “everything” applies to everything. Conventions and stipulations can’t create facts out of nothing: there must be antecedent facts that underwrite any convention that’s not empty air. If there is no fact—about Crusoe 3, about his dispositions, or about the world—that can explain why his word “everything” is to apply to *everything*, then it’s idle to claim that a convention he has adopted, or a way that he understands his words (apart from his dispositions), is doing this work instead. It’s equally “hot air” to demand of an “ideal interpreter” that he must idealize Crusoe 3’s words as picking out natural kinds.⁴⁷

Interim Conclusion: There is no way to make sense of the idea that

Crusoe 3 can make *any* word—no matter how fundamental—pick out a something, or a kind of somethings, in the world by his sheerly intending it to do so, and by (somehow) letting the world do the rest of the work needed. There is no way to even make sense of the idea that Crusoe 3 can *think* of doing this. It should also be clear that none of these conclusions depend on Crusoe 3’s isolation. A *community* of people, all busily pulling up on their own bootstraps, or on the bootstraps of one another, are no more likely to rise into the air, than an isolated Crusoe 3 so pulling is. Finally, it won’t help to introduce additional *ad hoc* metaphysical devices—such as “properties,” or Fregean “concepts,” or other sorts of *abstracta*—that Crusoe 3 is to supposedly “grasp,” and associate with his words, and that do the needed referential work for him. Why on earth should it be that he “grasps” just the kinds of prophylactic devices needed, and not other ones that fail to “match” the world in the contours they impose on Crusoe 3’s words?

11. A ROBINSON CRUSOE WITH CONTINUOUSLY VARYING DISPOSITIONS THAT HE HAS CONSCIOUS ACCESS TO

Let’s now consider Crusoe 4. He lives in the same kind of artificial world that our previous Crusoes have lived in. For purposes of exposition, however, let’s distinguish between Crusoe 4’s *current* dispositions to apply a word—at any moment—and his *previous* dispositions to do so. Imagine that Crusoe 4’s memory is very good, and so he can remember both the specific items that he applied his words to in the past, and the dispositions that he did this with. And, unlike Crusoe 2, who has two static sets of dispositions, Crusoe 4’s dispositions change over time. From God’s point of view, they exhibit a monotonically positive “learning curve.” That is to say, his dispositions to apply his words steadily and monotonically improve over time—given God’s standard of how well they match the natural kinds of the world. From God’s point of view Crusoe 4 gradually, but continually, gets better at identifying, distinguishing, and making similarity judgments about the objects in his world.

This is how *God* can put it. What *Crusoe 4* notices, instead, isn’t a monotonically positive *learning* curve, but rather that his current dispositions always enable him equal or more success in getting around

his world than any of his previous dispositions did. Because of this, he has adopted a systematic policy—when there is a disagreement—of always letting his current dispositions to apply a word trump his previous dispositions. In this way, Crusoe 4 has induced over time a monotonically positive *success* curve in his interactions with his environment. The same success curve is visible to Crusoe 4 with respect to his applications of his counting words: in the past he often found that he had to go foraging sooner than he expected; sometimes he found himself unexpectedly facing many more predators than he at first thought were there. But over time, as he adopts different methods of counting, these sorts of things happen less and less often. Sometimes he can even see how he has induced his success curve to improve over the course of a day.

God can see how Crusoe 4 keeps inventing better, and then still better, counting algorithms—algorithms that from His point of view enable Crusoe 4 to get the right answers (or to get answers closer to the right ones) in successively larger and larger proportions of the counting tasks he undertakes. For example, whereas Crusoe 4 used to count the coconuts in a tree in any old way, now he first sizes up how the coconuts are grouped in the tree, and then counts the clumps of coconuts in the tree that he's discovered he's likely to overlook otherwise. His counting dispositions have changed over time.

God has no reason, of course, to treat Crusoe 4's words as referring other than as his (total) dispositions (at a time) incline him to apply them. Because of Crusoe 4's policy of always letting his current dispositions trump his previous dispositions in how he applies words, God therefore, thinks of Crusoe 4 as developing, over time, a continuously-changing *series* of differing disposition-meaning languages, where the words of the later languages fit reality better than the earlier ones did. (God can talk this way because He can *see* to what extent Crusoe 4's changing dispositions match the world.)

As the earlier Crusoes did, Crusoe 4 has named collections of objects that strike him as similar. But he's not in a position to easily compare differing disposition-meaning *languages* in terms of their success, as Crusoe 2 did, because Crusoe 4's dispositions are constantly changing. Crusoe 2 individuated his (meaningful) words in terms of the differing dispositions associated with them. Crusoe 4, instead, thinks of

his words as *floating*—over time—in what they refer to because what they refer to-at a time-is given by his dispositions at that time, and that can change. If he thought to ask himself: “Am I still using these words as I earlier intended?” his answer would be: “Yes I am because I'm always resolved to use them as my current dispositions incline me to use them.” So Crusoe 4 can also understand his language to be a disposition-meaning language.

But, although Crusoe 4 can distinguish his current dispositions from earlier dispositions, as noted, he can't individuate the words by saying (as Crusoe 2 could) that: “I used the wrong word ‘coconut’.” He *can* say “I used the word ‘coconut’ with the wrong dispositions,” because he doesn't allow that his earlier *dispositions* to use his words indicate what his (current uses of his) words refer to. In saying this, Crusoe 4 acknowledges that his current use of “coconut” refers to what his current dispositions impel it to mean, and he's disavowing his earlier uses of the word along with his earlier dispositions.

However, Crusoe 4 realizes that the dictates of his current dispositions will be discarded by him tomorrow if his future dispositions then disagree with his current dispositions. How is he to say this? Perhaps by saying, “‘Coconut’ refers to this item today, but it might not tomorrow because I might hit on better dispositions to associate with the word than I have now.” I won't pause to analyze in detail how this usage of Crusoe 4's is supposed to work. In any case, something like this is reasonable for him to want to say (to himself) because he recognizes that the references of his words will continue to float in the future: they will always refer at the moment of utterance to whatever his current dispositions dictate them as referring to.

Crucial to the coherence of Crusoe 4's private-language practice is the monotonically-positive *success* curve that he can induce in his interactions with the environment by *changing* his dispositions in one way, and not in another. Notice that, as I've imagined Crusoe 4, he doesn't think that the range of his words “match up” with the way natural kinds are really configured; relatedly, he doesn't try to provide an explanation for his success in metaphysical terms. He has only adopted the policy of taking his words to mean what his current dispositions incline him to take them to mean because that's the best way for him to induce a monotonically-positive *success* curve. The meaning of Crusoe

4's words can be understood, therefore, as genuinely being guided by his intentions. After all, his ability to remember his previous dispositions to use his words, and that he isn't forced to obey the dictates of his current dispositions, is why he can be understood as having made a rational choice.

Two important points should be made now about the preceding discussions of the Crusoe cases. The first is this. I've placed my description of the interactional successes of the various Crusoes against a very neat and simple background metaphysics. I've even given an example of where Crusoe 4's monotonically-positive success curve corresponds to what we can call a monotonically-positive *learning* curve because, metaphysically speaking, his series of disposition-meaning languages have successive words that correspond better and better with how things group into natural kinds in reality.

It's important to realize, however, that such background metaphysical assumptions aren't crucial to the *coherence* of the various private-language practices that these Crusoes have. In describing the various Crusoes' private-language practices as "coherent," I mean two things: First, that such private-language practices contribute positively to a Crusoe's well-being. And second, that, where needed, a Crusoe has ways of evaluating better and worse alternatives among those private language practices.⁴⁸ What's crucial to the coherence of a Crusoe's private-language practice is that he has the capacity to evaluate the success of alternative approaches when he engages with his world, so that he can induce a monotonically-positive success curve.

The reason that the background metaphysics I've presumed in these Crusoe thought experiments isn't a necessary component of them is because a Crusoe can experience a monotonically-positive success curve for *all sorts* of metaphysical reasons. It needn't be, for example, that a Crusoe's successive series of languages are (from God's point of view) asymptotically approximating the natural kinds in the world. Clearly, *some* sort of positive dialectical interaction, between the ways the world is and the ways that a Crusoe can change his dispositions, is needed in order for that Crusoe to induce a monotonically-positive success curve; but this hardly has to involve the kind of simple and clean metaphysical interactions that I've hitherto presupposed. Perhaps, for example, there are no metaphysical kinds at all, and perhaps a Crusoe is pro-

tected from this interfering with his inducing a monotonically-positive success curve because of the strictly local and restricted nature of his interactions with the things in his world.

This point is especially important to stress because—as I argued in section 10—a Crusoe can't judge success by metaphysical comparisons of any sort. What this, in turn, shows is that a Crusoe's being able to induce a monotonically-positive success curve is itself necessary and sufficient⁴⁹ for his private-language practice being coherent. In what follows, therefore, when I assume that a Crusoe has a successful private-language practice, I'm (usually) just assuming that he has figured out ways to induce a monotonically-positive success curve in the series of private languages that he invents.

The second point is a caveat to my first point. I've described Crusoe 4 as enjoying a monotonically-positive success curve with respect to his interactions with his environment. But such a curve being monotonically positive isn't necessary for the coherence of a Crusoe's private-language practice. Crusoe 4 can take a wrong turn (for a while), or have a string of bad luck. He can go through a period where his current judgment is disturbed, he realizes that, and so he sticks (however well he can) to the dictates of his earlier superior dispositions.

Consider, for example, a Crusoe with a combination of the dispositions of Crusoe 2 and Crusoe 4. This Crusoe's success curve can shift in a negative direction, when he's tired for example, or drunk. As long as he has *some* idea of when he's tired or of when he's "off his game," his private-language practice will still be coherent. For he can let the meanings of his words be dictated by the best of his dispositions: sometimes, that is, this Crusoe lets his current dispositions dictate what his words mean, and sometimes he lets the dispositions that he had on other days, or earlier in the day, or tomorrow when he's well-rested, do so. The crucial element for his private-language practice being coherent is that he can identify better and worse dispositions in terms of their impact on his success curve, or (more indirectly) in terms of psychological signs (e.g., feeling tired or dizzy) that track more or less well the impact of those dispositions on his success curve.

It's an empirical question exactly what sorts of dispositions (and changes that can be induced in those dispositions) are compatible with a coherent private-language practice, and which ones aren't. If a Cru-

soe's powers are precipitously declining, because of Alzheimer's disease, a coherent private-language practice won't be possible. On the other hand, it could just be that a Crusoe's dispositions operate in an arbitrary way that he can't get a grip on at all. Here too, a coherent private-language practice is just impossible. I'm not (nor should I even try to be) prepared to describe specific conditions on the developmental processes of the dispositions of isolated individuals that allow a coherent private-language practice. One reason for this is that what such are isn't just a matter of the individual but of him *and* his environment. I'll simply describe any set of (changing) dispositions of a Crusoe (in an implicitly-given background world) as a set of "private-language-practice coherence-inducing" dispositions (or, a set of *plpci* dispositions) if they are due to a positive success curve that the Crusoe recognizes and induces. As just noted, whether an individual has such a set of *plpci* dispositions or not turns not only on him but also on what kind of world he is living in. And, as we've seen, the kind of *world* he's living in isn't something the Crusoe in question can independently assess by comparing it to his dispositions. He can only recognize that the changes in his dispositions that he's induced contribute to an overall positive success curve in his interactions with that world.

12. A ROBINSON CRUSOE WITH PLPCI DISPOSITIONS THAT HE HAS-AT BEST-INFERENTIAL ACCESS TO

Crusoe 2 and Crusoe 4 are psychologically unrealistic because I've presumed them to have perfect introspective access to their own dispositions to use words and apply concepts. It's hard to determine what degree of access to one's own dispositions to use words and apply concepts is psychologically realistic (to what degree, with respect to what sorts of dispositions, and how). I'd like, however, to now consider a Crusoe 5 with—largely—no introspective access to these dispositions: his dispositions to apply most words and concepts are entirely subpersonal. This is *not* to claim that all of his psychological states are altogether introspectively unavailable. That—in my view—*would be* psychologically unrealistic. Crusoe 5 has access to those psychological states (and so, he has words, and concepts—corresponding to these—for which he has access to his dispositions to enable his application of

such) that—as described in section 9—are needed for him to track his success. Let's also imagine that, given the world that he lives in, Crusoe 5 has a set of *plpci* dispositions. Because he has only indirect access to his own dispositions to apply (most) words, I want to describe a very different way that he must think of the meanings of the words of his private language—a way (I claim) that's quite similar to how we think about the meanings of the words of *our* (public) language.

Let's start this way. Since Crusoe 5 is unaware of his dispositions to apply (most) words, he's largely only conscious of changes in objects in the world. Objects in the world are sometimes one way to him and sometimes another. I can put it this way: Sometimes they *present* themselves one way to him, and sometimes another; sometimes they present themselves to Crusoe 5 as having certain powers, and sometimes they present themselves as having other powers. Sometimes, when he relies on how objects have presented themselves to him, he doesn't enjoy success; other times he does.

Because Crusoe 5 hasn't introspective access to his own dispositions to apply (most) words, he must recognize indirectly (or infer) any changes in his dispositions to apply those words by how changes in his dispositions have affected his interactions with objects. As already stressed, Crusoe *does* have introspective access to certain psychological *states* he can be in. Apart from recognizing that he is in pain, or afraid, and so on, Crusoe 5 can also recognize that he has become tired, or drunk, or dizzy. And he is capable of inferring that his dispositions to apply words have changed because he is tired, or drunk, dizzy, or in pain. But—as my use of the word "inferring" suggests—he recognizes this *indirectly*, because he has learned from previous experience that when he is in such states things present themselves in certain typically different ways from how they otherwise present themselves to him. Things are fuzzy, for example, or more difficult to handle, and so on.

A neat and effective way that Crusoe 5 uses to categorize the ways that objects in the world present themselves to him is to distinguish between what an item *looks* like and what it *is* like. The first question to ask about this distinction is *how* Crusoe 5 applies it to *objects*. Recall the discussion of section 10. There the point was made that a Crusoe cannot judge the adequacy of his words by comparing them to the world. The same point comes into play here. The distinction Crusoe 5

makes between how objects are and how they seem to be cannot be a distinction he draws by comparing how things appear to him to how they are.

How, then, *does* he draw this distinction? Let's first notice that he first learns that objects can shift in all sorts of ways. They are fuzzy, when he presses on the sides of his eyes; they have different colors in different light; and so on. The next question is how he can sort the various ways that he experiences how objects can be into "how they only appear" and "how they are." The full story, of course, is extremely complicated⁵⁰; but my purpose here isn't to even try to give a full and accurate story. My only aim is to give an *illustrative* story that shows that there is a way that Crusoe 5's verbal distinctions, between how things appear and how they are, his concomitant notion that he can be mistaken about how they are (because they appear differently from how they are), and his practice of correcting himself (later) in terms of this distinction, can be coherent.

My illustrative answer should not surprise anyone who has read this far: as noted, Crusoe 5 has plpci dispositions. And, generally, he describes how an object *is*—as opposed to how it appears—in terms of how he "corrects" himself later, when he sides with his currently better dispositions. By his "currently better" dispositions, I don't necessarily mean the dispositions he has *now* (for he may *now* be tired, recognize it, and therefore distrust the way things seem to him to be now—that is, he may distrust the dictates of his current dispositions). I mean, rather, those dispositions (at the current time) that belong to his set of plpci dispositions—the ones that induce a positive success curve.

An important point to make is that Crusoe 5's distinction between how things appear and how they actually are isn't a systematic one that he can learn to draw all at once. It's a distinction that Crusoe 5 takes himself to be fallible about. *Any* way that he currently thinks objects to be is one that he knows he can subsequently recognize he is mistaken about; so too, *any* way that he thinks objects (only) appear to be may be a way that he can subsequently also recognize that he is mistaken about.

As I've been indicating, this appearance/reality distinction is the basis of Crusoe 5's practice of describing himself as having made a mistake. His subsequent corrections of himself are accompanied by

his describing his earlier pronouncements as "mistaken," in comparison to his current pronouncements. And, of course, the coherence of Crusoe 5's so distinguishing between his making a mistake, and—more importantly—his correcting such mistakes, turns on his dispositions being plpci ones.

This isn't the place to dwell on these points in detail, or to refine them, so that the ways of speaking and thinking about mistakes that I've attributed to Crusoe 5 are more accurate to the appearance/reality distinction as *we* actually draw it. Given my description, however, a characterization of Crusoe 5's application of the appearance/reality distinction (against a background of a mistake/correction practice) to the meanings of his words is easy. He takes himself to be applying words to collections of things—to be naming kinds of things. Given that he already accepts that something can appear to be different than it is (and that he can subsequently correct himself on this), it easily follows that he can think it's possible for him to believe that something is a coconut because it *appears* to be a coconut, but actually isn't. In correcting himself, he takes himself to have now discovered the object to be a squirrel instead. He now *knows* the object is a squirrel although it earlier appeared to him to be a coconut.

Crucial, I claim, to Crusoe 5 being able to apply this kind of talk, about how things appear and how they are, to the meanings of his words is that he is only indirectly aware of his dispositions. He would not even *come up* with this kind of talk if he were directly (introspectively) aware of his dispositions all the time—that he was disposed to have such and such experiences when interacting with the world in this way or in that way. Instead he would be focused (as Crusoe 2 and Crusoe 4 are) on which changes in his dispositions induce the steepest positive success curve, and which don't. This isn't an option for Crusoe 5 because he isn't aware of his dispositions "in action," as it were. He can only think about them afterwards—in a more "philosophical frame of mind"—as what he theoretically posits as among the causes of why he thinks something is a coconut, when he does so. In a more searching vein of thought, he can recognize that the ways he categorizes things can only be due to his "dispositions," and not to "the ways that the world is." At that point he will be on his way to discovering Kripke's Wittgensteinian rule-following paradox.

Let's stay focused, however, on a less philosophically-minded Crusoe 5. As Crusoe 5 thinks about his own practices, he sees himself as using words to group collections of objects by their similarities; he thinks that he sometimes gets these similarities right, and that sometimes he gets them wrong. Furthermore, instead of thinking that his success is leading him to group items differently than he did earlier (to ensure that his success continues), he "puts the cart before the horse": he instead thinks that his success turns on his grouping objects correctly—as *they are in the world*. He's consequently prone to speak (to himself) this way: "I better get it right that the item up there is a coconut before I climb this tree" He thinks of himself as sometimes wrong, and sometimes right, *about whether something is a coconut*.

Because he speaks of his making "mistakes" in this way, it's natural for Crusoe 5 to take his invented words, "coconut," "1," "2," etc., to pick out various kinds of objects, and various cardinalities of collections of objects. When he hasn't made any mistakes, that's when he can pick these things out correctly, and that's when he correctly identifies what these words actually refer to. He, of course, thinks of his dispositions to use words—when he thinks of his dispositions at all—as entirely *defeasible*: at any time, he might get it wrong. He also thinks he has some idea of when he's worse at recognizing coconuts (or numbers of coconuts) than he is at other times; and he notices that he gets better at these tasks when he practices.

From God's point of view, what Crusoe 5's "practicing" comes to is this. Crusoe 5 changes his ways of going about his tasks—sometimes in small ways, sometimes in large ways—and he recognizes (sometimes consciously and sometimes subpersonally) when these changes lead to a positive change in his success curve. This is how (again from God's point of view) he is able to develop a set of plpci dispositions over time. That is, most of the time, as he experiments with modifications in how he engages in various tasks, he has more than one option in the directions he can develop his dispositions. If he hits on developing them in the right way, the resulting changes in his dispositions will exhibit a positive movement in his success curve that he can recognize; otherwise they won't. When I describe him as having several options, I mean something like this: He can count carefully or quickly; he can count the items in small groups, and total them, or he can count them

in any which way. He can do the task carefully and slowly, or he can rush it. And so on. Thus "what Crusoe does" must be understood in a very fine-grained way—what's relevant is that the changes in his behavior affect his answers. And, what causes Crusoe 5 to shift in what he does is whether a positive success curve is induced by one option that's steeper (in some ways) than the positive success curve induced by other options.

I've described all this from God's point of view. Crusoe 5's point of view is different: he's getting better at recognizing various kinds of objects, and he's getting better at undertaking various tasks: counting coconuts, for example. As he practices, he makes fewer "mistakes."

Notice that Crusoe 5's use of normative expressions ("correcting his mistakes," "being right," and "doing a task correctly") is coherent, and coheres with his private-language practice only because he has plpci dispositions. He corrects himself, and unbeknownst to him, it's useful for him to do this *because* the pattern of his corrections (the changes in his current dispositions) enables a positive success curve.

Crusoe 5, as I've indicated all along, labors under a serious misconception about how he applies his own words to items in the world. According to God, his words and their references float over time much as Crusoe 4's did, although Crusoe 5 doesn't recognize this. Instead, he thinks his words are fixed in what they refer to, and he thinks his dispositions (at least in principle) can always mislead him. As a result, as I've also indicated, Crusoe 5's language, and his beliefs about the words of his private language, are pretty much the same as ours (about our public language). Of course, Crusoe 5 thinks he *understands* what his words mean. He thinks, that is, that he can recognize collections of things that are alike (despite his tendencies to mistakes), and he thinks that he knows what functions his arithmetical inventions pick out. Where he has hit on rules for computation, he thinks he is following those rules when making calculations. He thinks these functions are *not* defined in terms of his dispositions (when he thinks of himself as having dispositions), and he justifies himself in his belief that he is applying the words as he intended earlier by his confidence in his abilities—always subject to correction by himself at later times. And indeed, he certainly can correct himself (later) when he (earlier) has had too much to drink; he may even, if he tries to calculate during

a drinking bout, say to himself: “None of this is right, probably. I’ll straighten it out after I’ve slept it off.”

As I’ve indicated earlier, Crusoe 5 is likely to remain complacent about these beliefs about his private-language practice unless his thinking takes an unusual (a philosophical) turn. Only then he is apt to puzzle himself with questions like: “How do I know the difference between *really* going on in the same way, and *only thinking* that I am?” and “What can it possibly be that determines what ‘coconut’ refers to, that’s independent of my dispositions to apply ‘coconut’?”

Let me stress one last time an important reversal in perspective that I’ve introduced in this section, and that’s crucial both to the coherence of Crusoe 5’s private-language practice, and to his nevertheless being mistaken about how that private-language practice actually works. Crusoe 5 thinks he is successful *because* he is getting better at recognizing which objects really are similar to which. But this has the situation exactly *backwards*: What’s actually the case is that his dispositions to group things as similar and different is changing over time—although its doing so is largely subpersonal. And what’s guiding this process—and making it not an arbitrary matter—is that the changes in his dispositions, that he describes as corrections of his earlier mistakes, are induced because they *lead* to greater success in his interactions with the items in his world.

13. FURTHER DISCUSSION OF CRUSOE 5: IS IT IN TERMS OF ASSERTABILITY CONDITIONS OR IN TERMS OF TRUTH CONDITIONS THAT CRUSOE 5’S PRIVATE LANGUAGE SHOULD BE ANALYZED?

As we’ve seen, Crusoe 5 has found the following kind of talk useful. He speaks of there being “tigers,” “coconuts,” and so on. That is, he speaks of his word “tiger” as referring to tigers, and his word “coconut” as referring to coconuts. Furthermore, he makes claims (to himself) about these things, and he regards some of these claims as true and others as false. He takes it that he can be wrong about any claim he makes. He thinks (he is currently disposed to think) that there are no tigers over there. But he thinks it’s possible for him to be wrong: It might be *true* that there are tigers over there. Some of this talk is similar to how the earlier Crusoes spoke; but some of it is quite

different.

Crusoe 5 talks (to himself) just the way *we* talk to one another (and sometimes to ourselves). Just like us, Crusoe 5 talks about what’s true and what’s false, he talks about his words referring to things in the world, and he takes it that he can be wrong about any of this. He can be wrong about *that* being a coconut, or he can be wrong about coconuts all being edible. Because he thinks he can be wrong, he thinks that the standards for whether he is right (or wrong) about whether something is a coconut or not is whether or not the thing in question *really is* a coconut.

Crusoe 5 takes his word “coconut” to refer to all and only the coconuts. He takes his word “five” to refer to all and only collections of five objects. In addition, he uses words like “true” and “false” of propositions he currently believes and of ones he once believed. He talks about correctly seeing that something is a coconut, and sometimes incorrectly thinking something was a coconut. He talks about making mistakes when he counts coconuts, and correcting those mistakes later. It’s the positive success curve, that changes in his dispositions induce, that makes this practice coherent.

God’s picture of Crusoe 5’s private-language practice is (unsurprisingly) very different. For God, as noted, Crusoe 5 doesn’t have one language; he has a continually-changing series of disposition-meaning languages. Or, equivalently, Crusoe 5’s words keep changing in what they refer to. (The *references* of Crusoe 5’s words *float* over time.) From God’s point of view, Crusoe 5’s words refer (at a time) only as Crusoe 5’s dispositions (at that time) dictate them to refer. But the extensions of Crusoe 5’s words keep changing, and in a positive way. Given *one* possible background metaphysics, God can see why Crusoe 5’s success in his world keeps increasing: the extensions of his words more and more approximate the natural kinds that there are. Given a different background metaphysics, God sees Crusoe 5’s success more ironically: due to an epistemic *cul de sac*, Crusoe 5 is enjoying what God recognizes to be a strictly local success.

Imagine a case where Crusoe 5 has dispositions that—at first—dictate no answer to the number of coconuts in a particular tree. There are simply too many there for him to manage counting consistently. When he attempts a count, he is disposed to recognize that his answers

are varying too much, and so he doesn't trust his abilities. Later, his dispositions change, and he can successfully produce an answer. He thinks: only *now* do I know that the number of coconuts in that tree is 57. It always was 57, of course, but I didn't know that until now. It was always *true* that the number of coconuts in that tree was 57 even though I was once incapable of finding that answer out. Perhaps there are trees out there with so many coconuts that it's impossible for me to ever know how many they have. Nevertheless, for each coconut tree, there is a number such that it is a fact that the coconut tree in question has that number of coconuts.

Was there a fact of the matter that there were 57 coconuts in that tree before Crusoe 5 had the capacity to give an answer? In asking this, despite appearances, I'm asking a question *within* Crusoe 5's language, using his word "57" and not my word "57." Putting sentences and words of Crusoe 5's language in c-quotes (to indicate that I'm quoting his language despite its homonymity to mine) resolves any danger of misinterpreting my question, for my question can now be posed this way: Did a fact correspond to ^cthere are 57 coconuts in that tree^c?⁵¹ One temptation is the old one discussed in section 3: to describe the facts in question as ones about how Crusoe 5's dispositions *will* develop (under such and such circumstances). Since he will (some-day) have dispositions that dictate ^c57^c as the answer to the number of coconuts in the tree, there is a fact of the matter about this. Crusoe 5, as we've seen, thinks that there is a definite number of coconuts in *any* coconut tree, even ones that it's impossible for him to *ever* be disposed to count. So, regardless of whether we give in to the temptation to characterize the facts about the number of coconuts in possible coconut trees in terms of how Crusoe 5 will be disposed to count, Crusoe 5 clearly takes his words to refer beyond the scope of his dispositions, and he takes his talk of what's true and false to extend beyond the scope of his dispositions as well. For us (and for God), there are no grounding *facts* that underwrite these claims of Crusoe 5's; consequently, there are no correspondence facts that determine their truth.

We (and God) might be tempted to translate Crusoe 5's terms to ours (His). In particular, we might think it appropriate to translate Crusoe 5's terms ^cnumber^c to our "number." In this way, we would legitimate the truth value for Crusoe 5's ^cthere are 57 coconuts in that

tree^c, and other statements of the same form, regardless of whether Crusoe 5's dispositions can *ever* provide grounding facts for them. *Doing so faces the original rule-following paradox.* There are no grounds that can be used to legitimate a translation to our "number" as opposed to an alternative "number*." Furthermore, it's hard to see what (a priori) constraints one could (legitimately) place on translation practices that would justify ourselves (or God) in so translating Crusoe 5's words. If the original rule-following paradox—as described in the opening sections of this paper—has force, then it has force here as well.

Someone who grants the coherence of the private-language practices of the various Crusoes might argue that what I've shown, rather surprisingly, is that assertability conditions don't require a community, as Kripke requires them to. Given that an isolated Crusoe 5 induces a plpci set of dispositions in himself (whether he realizes he is doing that or not), he is capable of a private-language practice governed by assertability conditions. Suppose, for example, that such a Crusoe 5 has developed dispositions to add. One might try this as a suitable assertability condition: Crusoe 5 is entitled to claim that he means addition by 'plus' because (roughly) he is confident that he can give correct answers to new cases, subject to correction by his better self—and where such a better self is characterized in relation to the Crusoe 5's plpci dispositions.⁵²

I *don't* want to claim this. Instead, I claim that Crusoe 5's language can be regarded as a "truth-conditional language." What's needed to support this claim is the denial of an absolutely critical assumption of Kripke's Wittgenstein—this assumption, of course, isn't restricted to him, but is widely-held—that truth conditions *require* correspondence facts. As various sorts of truth-deflationists have argued,⁵³ however, the sentences of a language can have "truth conditions" even if truth isn't a correspondence notion.⁵⁴

I'm going to develop this point in a slightly different way—a way that I've put the matter in other work. It can be *indispensable* to a language practice—private or public—to describe sentences as either true or false, even though some of those sentences have terms that don't refer.

Consider, for example,

- 1) Neither Mickey Mouse nor Minnie Mouse has *ever* been depicted in movies as a plumber.

This sentence is either true or false. But it is not so because there are facts about *Mickey Mouse* and *Minnie Mouse* that (1) corresponds to, and that it either correctly or incorrectly describes.

Crusoe 5's way of talking may seem to apply truth and falsity to sentences of a more radically indeterminate nature. There is, after all, something about the world that's relevant to an evaluation of the truth value of (1); one can even operationalize a procedure for determining its truth. But consider a claim about the number of items in some collection that (by assumption) outstrips Crusoe 5's dispositional resources now and forever. Is it supposed to be true even in this case that there is some number that is the number of items in that collection?

Yes. Crusoe 5 speaks (as we similarly do) of not 'knowing' whether something is the case or not. Crusoe 5 takes it that every tree containing coconuts *has* a particular number of coconuts. This follows from his understanding that "coconut" refers to all and only the coconuts independently of his inclinations to so apply the word. He has the same understanding of the applications of cardinal-number words. Crusoe 5 doesn't recognize that the references of his words float to the dictates of his best dispositions at a time; and so he has a complementary way of speaking of his ignorance, of his not knowing something, that makes no distinction between three kinds of cases: (i) where he doesn't 'know' something for which his dispositions (at a time) *will* provide an answer if he investigates, (ii) where he doesn't know something for which his dispositions (at some later indeterminate time in the future) will then be able to provide an answer, although none of his dispositions can do so now, and (iii) where he doesn't know something although his dispositions can *never* provide an answer.⁵⁵

Crusoe 5's words, 'true' and 'false', it must be stressed, are his words (as the c-quotes graphically indicate). And so (according to God), just as all his other words float to the dictates of his best dispositions, so do these words. Naturally corresponding to this practice with 'true' and 'false', as I've just noted, is a usage of ignorance, of not 'knowing', that doesn't distinguish between cases (i)-(iii).

Should God accept the cogency of Crusoe 5's use of the words 'true' and 'false'? He can if He wants to, by adopting Crusoe 5's language. His understanding of Crusoe 5's language can take this form: 'It's true that there are such-and-such numbers of stars or there isn't'. However,

although God knows of every statement in His own language which is true and which is false, this isn't so of Crusoe 5's language. In particular, God doesn't know how to fill in the blank in the sentence of Crusoe 5's language which reads: 'the number of stars is —'. He can't do this even if He translates 'star' into His own language as 'star', and He translates Crusoe 5's small numbers into His own (small) numbers. This is because he has no justification in translating Crusoe 5's cardinal words as his cardinal words *when these outstrip Crusoe 5's dispositions*.⁵⁶ God's omniscience is unlimited in His own language; but it's limited when He chooses to understand the words in Crusoe 5's language. His omniscience is limited by the fact that in Crusoe 5's language talk of what's true and what's false outstrips any grounding facts that can be used to determine the appropriate correspondences between statements and the world.

I should point out that these claims have—strictly speaking—nothing to do with the fact that God is allowing Himself to understand Crusoe 5's *language*, as opposed to Crusoe 5's *thoughts*. The same point holds even if Crusoe 5 doesn't have a language, but just thoughts and concepts. In this case, such thoughts and concepts are disposition-meaning thoughts and concepts, and so God recognizes that their scope and range goes just so far as Crusoe 5's dispositions allow them to. The absence of grounding facts for Crusoe 5's thoughts leads to the same result for God: When He chooses to understand Crusoe 5's thoughts, his omniscience is correspondingly limited: He does not know—with respect to many collections, 'how many there are'.⁵⁷

14. FLOATING STANDARDS. FEELING GUIDED BY ONE'S UNDERSTANDING OF WORDS

It has often been said that what makes private-language practices incoherent is that because a Crusoe is isolated, whatever is going to seem right to him *is* right (Wittgenstein §258). One reason this has been thought to follow from the sheer fact that a Crusoe is on his own, is that it has seemed that in order to deny that whatever is going to seem right to a Crusoe is right requires that there *be* a standard that a Crusoe's dispositions to use a word must be compared to, and that Crusoe himself must be able to compare his own dispositions to use a word to

that standard. More simply put: there must be an external standard of some sort, and this external standard must be one that a Crusoe can *use*.

But this *isn't* needed. *All* that's required is that, first, a Crusoe *understand* any of his applications of his words to be defeasible—that he believe it to be possible that, at any time, he might be “wrong” in his application of one of his own words. What's required, second, is that a Crusoe *have* a coherent practice of correcting himself over time. That is to say, what's required is that his practice of correcting himself not be arbitrary, erratic, or random, and further, that there be a genuine utility for a Crusoe to so correct himself (over time).⁵⁸

What I take myself to have shown, first, is that if a Crusoe has certain sorts of dispositions, certain abilities to change his dispositions, and a capacity to evaluate the impact of those changes on how well he can navigate his world, there will be a utility in his so-changing his dispositions in order to better navigate his world. Second, if his dispositions are introspectively inaccessible to him, he will be compelled to describe this practice, not in terms of “changing” his dispositions to better navigate his world, but rather in terms of “correcting” his false starts and wrong moves in how he categorizes the items in his world. He will instead have to describe himself as learning to better recognize the things in his world, and how they are similar and different from one another.

Although I've not dwelt on this, I want to suggest that the difference between feeling that one's dispositions to group things in the world in such and such ways *drive* or *compel* one's answers, and the very different impression that one's grasp of concepts of how things are grouped in the world *guides*⁵⁹ one's answers, is due entirely to whether one's dispositions are subpersonal or not. When they are subpersonal, one can only focus on items in the world, and how one appears to be correcting one's earlier impression of them. This is where the phenomenological impression that one is guided by one's concepts comes from. One would have a very different perspective if one were aware of one's dispositions, and how one was endeavoring to change *them* (or have *them* change) in light of whether doing so or not would induce a more positive success curve. It's in these ways only that Crusoe 4's impressions of the relationships between his dispositions, his words, and the world

differs from Crusoe 5's impressions of these relationships.

Third, I've indicated that these private-language practices of an isolated individual are compatible with a way of speaking of the truth and falsity of his statements that—with respect to most of his statements—outstrips the facts that can determine whether such are true or false. This way of speaking is both available to him, and it's available outside of his language—if it's translated into another, for example. Furthermore, his private-language sentences are compatible with truth-conditional approaches to his private language. This third claim, however (turning as it does on the cogency of a deflated notion of truth that doesn't rest on facts), is something I've not argued for here, although I've extensively argued for it elsewhere.⁶⁰

Here is an important corollary of this claim. A solution to the rule-following paradox that substitutes assertability conditions for truth conditions not only eliminates the cogency of private rule-following; for similar reasons, it eliminates the cogency of important aspects of our self-ascribed ability to *think*.⁶¹ Our thoughts can only be cogent in a setting where they correspond to sentences or other public vehicles of communication that we are entitled (subject to the ratification by others) to assert. What is implied, therefore, is a kind of *public sententialism*: thought is cogent only if it is normatively constrained by community standards. This corollary—I must stress—does *not* follow from my sceptical solution to the rule-following paradox simply because truth conditions are employed instead of assertability conditions.

One important point I've tried to make in this paper is this. Many philosophers have thought to draw a sociological conclusion from the rule-following paradox. On my reading, this mistake is a symptom of an overly-narrow appreciation of what's required of a coherent normative practice of correcting oneself and others—that a coherent practice of correcting “errors” requires a *fixed* standard that's external to the individual being corrected. This forces the view that such an external standard must be located in the community to which that individual belongs. And *this* motivates what I've called the straight sociological solution to the paradox. The other approach—Kripke's Wittgensteinian one—more radically replaces truth with public assertability conditions. Because of the nature of such assertability conditions, practices of correction can only be cogently applied in a public setting. Thus a com-

munity standard is also required by that view.

When we think about our ordinary practices, however, we realize that there is a tension between them and *any* sociological solution—sceptical or otherwise. It's part of *our* practice to regard it as possible for *one person* to be right, and for everyone else to be wrong. We allow that the ability of one person to sort certain kinds of objects may turn on his having certain skills, and on his being able to detect differences in things that others aren't capable of seeing. We also allow that it's possible for it never to be recognized (by other people, or by the society at large) that such a person is "right." More strikingly, we allow that it's possible for an entire society to be wrong, and for that society to never be corrected.⁶² One way to imagine such a thing is to imagine a Friday whose counting practices enjoy a superior positive success curve compared to those of every other individual in that community, so that he's capable of exploiting anyone else in (for example) barter exchanges. (He can, that is, use the *entire society* as a money pump.)

My approach, despite being a sceptical solution to the rule-following paradox, can accommodate the possibility of either an entire society being wrong, or an entire society being wrong, and one individual (who exploits the practices of that society) being right, because it bases the cogency of someone correcting her own practices in terms of her subsequent enjoyment of an enhanced positive success curve in her interactions with the world. The cogency of someone being right or wrong, therefore, is based in the interaction of her developing dispositions with her environment, *regardless of whether she realizes this or not*.

It's worth reminding ourselves of the powerful intuition we all have that a Robinson Crusoe is capable of developing a coherent language (and accompanying concepts) that he can apply successfully to the objects in his world in conformity with his previous intentions, and yet one that's not to be understood as relative to the standards of any community. As I've already stressed, it's these intuitions that make the "private language argument" so shocking in its import. Indeed, we all have the powerful intuition that our (own) grasp of counting enables us to understand how to go on *ad infinitum*.⁶³

My approach, being a sceptical one, ratifies these intuitions only in the very weak sense that it's *true* that we (and the various Robinson Crusoes) understand how to go on. There is no matter of fact about

this understanding—about our capacities—that fully underwrites this truth. What this, in turn, shows is that our ordinary talk of "understanding" (since it belongs to our language) is just like all the other words in our language in going beyond "the facts," what, in this case, any study of us—psychologically and subpersonally—can ever reveal. "Understanding," is an ordinary word that we use to indicate confident competence at certain tasks and in certain ways. It should be no surprise, therefore, that like "coconut," like the number words, and like pretty much all our words, the standards for it also *float*.

In the case of "understanding," the way the word floats to what it applies to at a time involves (in the case of counting, for example) later, and better positive-success-curve inducing, dispositions due to people learning first how to count with their fingers, and then to count by means of various tools (calculating devices: counting boards, etc.; notational devices: Arabic notation, etc.)—eventuating in the use of sophisticated devices of all sorts. At each stage in our acquisition of new approaches to counting, we gradually iron out various "mistakes": various earlier methods of counting that induce a less optimal positive success curve. We don't think of it *this way*, of course. We think instead that we are continuing to count in the same way we always intended to count. And (factlessly), *we are*.

15. IS OUR CONCEPT OF A GOD'S EYE VIEW COHERENT?

Here's a worry some philosophers may raise.⁶⁴ I've apparently introduced a standard of objectivity via a God's eye view of how our words should be applied that contrasts with how we actually apply our words. I've also suggested—in sections 8 and 10—that Crusoe 2, Crusoe 3, and Friday don't have access to God's perspective. More dramatically, they apparently can't even make sense of certain crucial notions that seem to be needed to characterize a God's eye view: they can't make sense of words better fitting (or worse fitting) the world. They can, of course, contrast one language with another by the varying positive success curves they induce; and they can fault the dictates of one language in terms of another by translating the terms of one language to that other, and treating the deviations as mistakes.

This point can be put—and it has been put by many philosophers—

this way: No one, including us, can escape their own language frameworks, conceptual schemes, or whatever. But when the point is put this way, then it seems to imply that Crusoe 5 (and ourselves) also can't make sense of the idea of a language fitting better or fitting worse to the world and (correspondingly, therefore) neither Crusoe 5 nor ourselves can make sense of a God's eye view of the world either. Thus—the objection goes—a God's eye view has been shown to be incoherent. Any contrast between such a perspective and our own is, as a consequence, nullified since we can't make sense of God's end of the contrast. Therefore, the best sense of anyone's words being right or wrong can only be in terms of how we use our words in our (current) language-framework, and our ability (which the Friday of section 8 shares) of faulting—relative to a translation—the words of others when they deviate from our own words.

I think, however, that the matter is more complicated than the above makes it sound. This section is dedicated to probing this surprisingly delicate question.

I'll start with the point that our picture of how the words of our own languages operate is largely the same as that of Crusoe 5. Let's consider the coconut again. Our term "coconut" refers to coconuts. We take this to be the case even though we treat our dispositions—even our collective dispositions—as all potentially untrustworthy. This is how we understand the relationship between our dispositions and (nearly) all our words.⁶⁵

One view—I'll call it *the contrastivist view*—is that this paper has shown that this picture is just *wrong*. The right view is that we (and Crusoe 5) utilize a continuously-changing series of disposition-meaning idiolects, where each of the words in each of these idiolects refers exactly to what the dispositions of the speakers of those idiolects (at those times) impel it to refer to. We (and Crusoe 5) suffer from the ineliminable illusion (because our dispositions to apply words are largely subpersonal) that the words in our language are largely stable in what they refer to.⁶⁶ We—but not Crusoe 5—think of our language as *public*, although we and Crusoe 5 agree that words are frozen in what they refer to insofar as they are frozen in what they mean. Due to great variations in knowledge (both among shared-language speakers, and over time for any particular language speaker), however, we vary greatly

in our abilities to apply those words correctly. "Coconut" is an example of a word that—we think—doesn't change in what it refers to over time; rather, we learn (over time) that some things that we thought were coconuts are in fact not coconuts (and vice versa). According to the contrastivist, that our words and their references float over time is something we misperceive as an epistemic fact: we think that our *knowledge* of the objects our words refer to is what floats over time and that the references of our words do not float.

Corresponding to these misconceptions is a correspondingly misguided semantic theory for our language. Such a semantic theory, for example, presupposes that a certain class of declarative statements is two-valued,⁶⁷ and it assigns truth conditions to those statements along standard Tarskian lines, based on what the words in those sentences refer to.

According to the contrastivist, the semantic reality is very different. Each person's plpci dispositions induce disposition-meanings (for that person) for each sentence S: such-and-such circumstances when each individual would assert S, deny S, or fail to respond in either way (e.g., by a "don't know" shrug, or by some other failure to assert or deny S). These disposition-meanings change over time because of the positive way that these changes enable individuals to succeed in their environments, and to communicate with one another. The nature of such changes is what (i) causes those individuals to enjoy the illusions of language described in the previous two paragraphs, and (ii) allows speakers of these changing idiolects to instead speak both of the truth and falsity of statements and of the truth conditions for such statements. In reality, however, statements have no truth conditions, at least not in the way that the misleading Tarskian semantic theory I alluded to takes there to be. Instead, truth and falsity conditions (this is one possibility) for each idiolect arise from the plpci dispositions of a speaker of that idiolect at the time that speaker possesses that idiolect.⁶⁸

It might be thought that I've too narrowly described the resources that plpci dispositions offer to semantic theories. We can take a chronologically longer view of them and characterize truth conditions only by the standards of later dispositions. Only with respect to the later dispositions, and only when they eventually stabilize on verdicts for certain

sentences, do we take those verdicts to indicate actual truth values. Doing so, however, still leaves numerous statements without truth values; furthermore, we do not normally think of a stable verdict on a statement as necessarily a correct verdict.⁶⁹

The contrastivist view that I've been describing in the past few paragraphs posits a sharp contrast between the reality of our languages and their relationships to our dispositions and the illusory picture of these relationships on the part of speakers that deeply misconceives them. The contrastivist view is one I've repeatedly attributed to God—but that doesn't make it *right*. There is another possibility, what I'll call *the noncontrastivist view*. The noncontrastivist denies that God's perspective of our languages falsifies our perspective of them. The relationship, rather, is like that between faces as they look to us normally, and what faces would look like if one could see the anatomy beneath the skin. The perspective in terms of plpci dispositions (the anatomy) describes what might be called the "engineering" of reference. But our words *do* refer to what they refer to, and part of the explanation of how they do so involves the trajectory over time of our (collective) plpci dispositions. The noncontrastivist view takes "coconut" to refer to coconuts. This is a reference relation between a kind term and objects in the world, one that is named by the words "refer," "reference," and so on. The "correspondence relation" is the relationship between the word "coconut" and certain objects in the world—coconuts.

Recall the distinction between grounding facts and correspondence facts that was given in section 1, footnote 5. The noncontrastivist view is that the correspondence relation is only *partially* instantiated by correspondence facts. The absence of (certain) correspondence facts arises from absences in the grounding facts in individuals (taken individually or collectively). For, first, as we've seen, the (individual or collective) dispositions to understand how to apply one's words in certain contexts may simply determine nothing at all about a great number of things—e.g., whether they are coconuts or not. Second, the dispositions may determine falsely (from a later perspective) that items are (or are not) coconuts.

In describing the reference relation this way, the noncontrastivist simply extends how we (and Crusoe 5) talk about the numbers of coconuts in trees, and other similar claims about what is or isn't so-and-so

(independently of whether we can ourselves determine it or not) to the very word "reference" itself—at least as *we* use it. Therefore (according to the noncontrastivist) it doesn't have to be denied either that "coconut" refers to coconuts, or that "refers to" characterizes the reference relation. Rather, the noncontrastivist is claiming that only *some* instances of the reference relation go beyond the facts.

To make the noncontrastivist view clearer, it's valuable to distinguish between the *applications* of a word and what it *refers* to. What a word applies to (at a time) is what the (collective) dispositions of its users (can) apply it to. On the noncontrastivist view, what *floats* isn't the reference of words, and it isn't (*pace* the contrastivist view) that words in languages are constantly superseded by successor words in successor languages. Rather, what floats is the applications of a word. Its reference remains stable (unless the word is discarded or officially changed in what it refers to).

On the noncontrastivist view neither reference floats nor is it that words are continually replaced by new words with differing extensions. But the standards by which these words are applied and, correspondingly, how mistakes are recognized isn't to be characterized in relation to what such words refer to but to the floating applications of these words. Therefore, just as on God's view, the *standards* by which our words are applied *float*. God has it right about how we apply our words and how we recognize ourselves to have made mistakes. But, the noncontrastivist claims, *pace* God, we should not conceptually connect our practice in implementing those standards with what the words refer to—in the sense that what the words *refer* to is the standard by which our words are to be judged as correctly or incorrectly applied.

What a word refers to we understand to be its application when there are no longer any mistakes being made. "Mistake," as here used, is similar to "not knowing," as characterized earlier: the words go beyond any description—now or ever—of how someone could modify their dispositions to yield a different (i.e., "correct") answer. (*All* our words, on the noncontrastivist view, are *partially* supported by the facts in what they refer to.) This negative characterization of the reference of a word, therefore, allows us to speak—for example—of there being a certain number of items (regardless of whether there are any grounding facts whatsoever that underwrite what that number is). Similarly,

we can speak of words referring to collections of objects regardless of whether there are any grounding facts that (fully) underwrite that reference relation. In general, there are various possibilities. The (collective) dispositions to apply a word, and “coconut” may be one of *these* words, floats for a time before the word stabilizes entirely in how it is applied. Our dispositions to apply the word no longer change although—of course—that isn’t true of our dispositions to apply other words. Furthermore, defeasibly, these dispositions actually characterize what the word refers to. Another case is where a word stabilizes in how it’s applied with respect to various items, and as our dispositions improve, it remains frozen on those items but continues to accrue new ones. Our ability to apply ever-larger cardinal numbers to ever-larger collections of objects fits this case. Again, this growing set of dispositions, defeasibly, actually characterizes larger and larger collections of items that the cardinal words refer to.

In general, if a word can be applied usefully at all, this is because our dispositions to apply it are stable on at least some class of objects, and it may become stable on a larger class of objects over time. However, there are always appropriate areas of the application of words that elude stability (and even application), and that will always do so. Nevertheless, just as we speak of the possibility of something being a coconut, or not regardless of whether our dispositions (now or ever) to use the word determine an answer, so too, we speak of words referring to objects independently of whether there are any grounding facts—now or ever—that underwrite that characterized reference relation.

I’ve elsewhere discussed what I call the “truth maker assumption”—the view that if a statement is true, then it is true because there are things (that the statement is about) that it is true of. (Recall the discussion of the sentence (1) in section 13.) There are many different versions of this assumption but they share the perspective that a statement that isn’t about anything can’t be true or false. If a statement *is* true (or false), then there is something that it is true (or false) *of*. This assumption is one that I think we should resist; and the noncontrastivist view reveals a way of resisting it that’s different from arguments against it that I’ve given in other work. There are no correspondence facts about whether “coconut” refers or not, with respect to many items in the world—for the grounding facts about how we use “coconut” don’t de-

termine such facts. *Nevertheless*, it is true (or false) of each such item whether “coconut” refer to it or not.

I should stress the point that the noncontrastivist view isn’t restricted to sentences and language: it extends to our view of our own thoughts and self-ascribed concepts. We understand our concept of *coconut* to pick out all and only the coconuts—even if we can’t (and even if there are no facts that determine what all and only the coconuts are). Built into our concepts, into our understanding of our concepts, that is, is that our concepts apply to what they apply to even though the grounding facts about us do not determine appropriate correspondence conditions (appropriate necessary and sufficient conditions) to sort everything in the world into what the concept applies to and what the concept doesn’t apply to. In the more ordinary language of section 2, we take our concepts to be ones we can grasp even if (as in the case of calculation) we’re not particularly good at successfully applying them.

I also want to stress again that the noncontrastivist view, as we have just seen with the word “refer,” extends the only partial grounding in facts of words (concepts) like “coconut,” “gold,” etc., to words and phrases like “reference,” “everything,” “mistake,” “know,” “understand,” and so on. These latter concepts, as well, can be ones that occur in sentences that are true despite there being no correspondence facts that they are correlated with.

Recall the issue that this section opened with. I gave the impression in sections 8 and 10 that Crusoe 2, Crusoe 3, and Friday don’t have access to God’s perspective. More dramatically, they can’t make sense of certain crucial notions that are needed to characterize a God’s eye view: they can’t make sense of words better *fitting* (or worse *fitting*) the world.

The noncontrastivist view creates conceptual room for the idea that our words refer to what they refer to *despite* our epistemically faulty dispositions. In exactly the same way, conceptual room has been created for the idea of a being whose dispositions aren’t epistemically faulty the way ours are: a being who simply knows enough so that He can apply His terms as they refer.⁷⁰ In so imagining His words suitably applying to the world, it isn’t that we are to imagine the world itself as shaped such and so (as it were) and His words as conforming to those

shapes. To do this really is to step outside our own conceptualizations and to try to compare His conceptualizations to the world. So too, if we do something similar with ourselves—imagine our words as fitting well or ill with the world’s own contours—we are illicitly stepping outside our own conceptualizations and trying to imagine (from outside those conceptualizations) a comparison of our own concepts with the world. What we *can* imagine is that the floating processes—that the applications to the world that our words go through as our dispositions are refined—can come to an end. And we can imagine that they have come to an end with respect to a word because the positive success curve induced by our current dispositions is already maximal; nothing will make it better. Similarly, we can imagine that God’s conceptualizations are already maximal in the positive success curve they induce with respect to *all* His words (and concepts). He can’t do anything to make them better.

Characterizing our notion of our words “suitably applying to the world” in terms of the content *maximal success* doesn’t involve an implicit picture of a world that we match our words to, nor does it allow the illicit thought that we can use this so-called *suitably applying to the world* as the standard by which we can take our words to refer.⁷¹ The standard for correcting our referential mistakes is always the same one: introducing a more positive success curve; and we always treat our applications of our words at any time as defeasible: they may not really refer to what we think they refer to because the positive success curve in question may not be maximal. Nevertheless, we can make sense of a being who need not use His words defeasibly because His practices with His concepts are already maximally successful.

Recall the discussion in section 8 where the question was raised of whether Crusoe 2 could make sense of a language better fitting the world. One theme—stressed there and throughout this paper—is that such a metaphysical notion can’t function as a standard for how words are used. The noncontrastivist view accepts this. In claiming that we can speak of what our words refer to—apart from how they are applied—it is not therefore being claimed that what they refer to can be used as a standard of correction for our words. The standard for the correct usage of our words is a floating one. The second important theme raised there is that it isn’t possible for Crusoe 2 (and

by implication, for any of us) to think of a language as “fitting” the world. I have pressed, instead, the idea that our notion of God’s application of His words needn’t be understood in this metaphysically loaded way. Instead, as noted, we can think of Him as epistemically faultless: as using his concepts (and words) in such a way that His positive success curve is maximal. We can accept the fact that there is no further content to our notion of words—ours or His or anyone else’s—“fitting the world” that goes beyond epistemic faultlessness: that the divisions the words in question introduce induce a maximally positive success curve.⁷² The worry was raised at the beginning of this section that the best sense anyone (us or Friday) can make of someone being wrong is by comparing their usages of words to our own—relative to a translation—and criticizing theirs for deviations. But we can now see that this is wrong. Thinking of “objectivity” as an epistemically faultless God’s eye view of things suffices to provide us with a notion that goes *beyond* our own current usage and dispositions to apply words. Thus this characterization of “objectivity” doesn’t amount to the empty ratification of how we currently apply our *own* words.

What has been shown in this section is that the question of whether or not we can make a kind of sense of the idea of an epistemically faultless application of our words—say by a epistemically faultless being such as God—comes down to whether the contrastivist or noncontrastivist view is the right one. On the contrastivist view, the idea of such a being arises from adherence to the wrong picture of our language. The right picture of our “language” is that we actually shift through a series of disposition-meaning languages over time, and not that we have one language that we apply the words of by our floating dispositions, along with floating standards of the correct application of those words. The noncontrastivist view opposes this by demoting the time-relative disposition-meanings to mere application-conditions of words at a time. The references of our words are instead characterized as those usages of them, if any, that would induce a maximally positive success curve.⁷³

As the foregoing discussion indicates, the contrastivist and noncontrastivist views are not terminologically disguised versions of one another. According to the contrastivist, we think we speak one language, but we speak a series of them over time: we are in the grip of an illu-

sion. According to the noncontrastivist, we speak one language just as we think we do. Furthermore, reference is stable, just as we think it is. If we are confused, the noncontrastivist claims, it is about the nature of the reference relation: we mistakenly think that it is always determined by facts when this is only partially the case (as with *all* words). This mistake about the reference relation, arguably, is one that only philosophers are in the grip of—not ordinary people.

Now consider the philosopher who wants to claim that the notion of a God's eye view of the world (a view of the world, that is, that isn't epistemically faulty) is incoherent. It seems to me that it has been shown that this claim faces three obstacles. The first obstacle is simply that the noncontrastivist view must be shown to be *wrong*, that is, it must be shown that what floats is the *references* of our words (or the words themselves) and not the mere dispositions to apply those words—that the reality is that we speak disposition-meaning languages. This opposes the noncontrastivist picture that the reference relation (like the relations of *all* our words to the world) is at least partially factless; and this, I want to suggest, is harder than it looks.

Given, however, that the noncontrastivist is wrong, two obstacles still remain. The first is this: to show that the idea of a God's eye view is *incoherent*—because it arises out of the wrong picture of our language—requires showing that in some strong conceptual sense a view of language that allows the possibility of a God's eye view is not only wrong but *impossible*. It's not enough to show that it's been empirically ruled out. I'm not at all sure that the contrastivist view, as a solution to the rule-following problem, *must* be taken to be a “conceptual analysis.” But whether this is so or not has to be sorted out before the *incoherence* charge can be entertained, let alone sustained.

The second obstacle is this. Suppose the perspective of language that the contrastivist treats as an illusion—that reference is stable and that our shifting relationship to it is an epistemic one—is one that it is impossible in practice to shed. I think that the discussion of Crusoe 5 indicates why this would be: the largely subpersonal operation of our and Crusoe 5's dispositions prevents either of us perceiving the references of our words to be constituted by our individual or collective dispositions. And suppose, for similar reasons, that it's also impossible to shed the same impression about our concepts. It then follows—because

of this illusion—that it's built into our *conception* of our concepts that there is an in-principle right way that they can be applied. One way to see this neatly is to notice how easy and automatic it is for us to *understand* that how we are applying our concepts can almost always turn out to be wrong. It's not clear how we are supposed to be able to *understand* this, given that the contrastivist view rules out an objective (epistemically faultless God's eye) view.

For given that we *do* understand our own application of our words and concepts to be fallible (apart from “pain,” and the like), it follows that we also understand how—lacking that fallibility—we would have a God's eye view of the application of these words and concepts. But then, how can it be that a God's eye view is a notion that we *don't* understand?

Notes

¹My thanks to Brendon Larver and David Corfield for organizing the conference *Two streams in the philosophy of mathematics: Rival conceptions of mathematical proof* (July 1-3, 2009, University of Hertfordshire, Hatfield, UK), and, of course, for inviting me to speak. My thanks to the scientific committee, and the invited organizers, of the 5th International Symposium of Cognition, Logic and Communication for organizing the conference *Meaning, understanding and knowledge* (August 7-9, 2009, Riga, Latvia), and of course, for inviting me to speak. My thanks as well to the participants of both conferences for making them such rich and enjoyable experiences. My thanks to Douglas Patterson for his detailed comments on the penultimate draft of this paper. They were extremely valuable, as always. (I'll probably regret not taking all his suggestions.) I should also add that I'm under the impression that we continue to disagree on a number of substantial points that come up in this article. (I'm hoping these disagreements make it into print sometime in the future.) My thanks as well to Eric Schliesser for directing me to some invaluable terminology. I should finally add that this paper discharges a rather old promise I made in my 1994, where I first indicated that my philosophical approach to mathematical practice had as a corollary a response to (a version of) the rule-following paradox. In spirit, although not in expository details, *this* is what I had in mind. The characterization of the evolution of scientific language that I subsequently presented in my 2000 (especially Part IV) illustrated the “floating picture” of language (given here); but reasons of space again prevented an explicit discussion of rule-following. I'm thus very grateful to the editors of the 5th issue of the *Baltic International Yearbook for Cognition, Logic and Communication* for the opportunity to finally lay all this out.

²All references in this article to Kripke's work will be to his 1982; all references to Wittgenstein's work will be to his 1958, except when otherwise specified.

³There is some literature on this question. See, e.g., McDowell (1984) or Goldfarb (1992).

⁴See, for good discussions of this, and the citation of relevant literature, Carey (2009)

and Butterworth (1999), chapter 3, sections 1-4. A classic study is Gelman & Galistel (1986).

⁵It's not particularly off-stage in Wittgenstein himself, because he often gives examples where people are counting objects. And Kripke does take note of collections being the target of counting when he discusses—in passing—“quounting,” and when he discusses Wittgenstein's text. However: if the counting of collections of *objects* is centrally incorporated into the examples that Kripke's sceptic challenges that will make salient possible (but overlooked) coherent forms of private-language practices, as I show in sections 7-12.

⁶So there are two uses of “fact” occurring here that should be kept distinct. The first are the purported (external-world) facts that true meaningful statements correspond to. The second are the psychological/dispositional facts in the person that underwrite her capacity to understand true meaningful statements, and consequently that determine that such statements are indeed understood to correspond to the (external-world) facts that they purportedly correspond to. Rule-following considerations place pressure on the second set of facts; the correspondence picture that requires the first set of facts is thus indirectly undermined. In what follows I'll describe the second set of facts as *grounding facts*, and the first set as *correspondence facts*.

⁷Some of the tacit knowledge required to do this successfully is labeled, by Gelman and Gallistel (1986, 77-82), as the counting principles: “cardinal word,” “order-irrelevance,” and “abstractness”—that the last word in a count is the cardinal number of the collection, that the order in which the objects in a collection are counted doesn't matter, and that it doesn't matter what the objects are. These abilities are acquired by children in stages from ages 2 to 4. (See Carey (2009, p. 241-244) for a description of the process, and for indications of how arduous this acquisition of tacit counting knowledge is.)

⁸See Butterworth (1999), especially p. 52-62, for an accessible discussion—with some citations—of the different types of number vocabulary to be found in different natural languages.

⁹It seems, however, that children can recognize the indefinite nature of numbers without necessarily having learned a genuine numeral system. Carey (2009, 252) quotes one five-year old as saying: “suppose you think a gazillion is the highest number—well, you can go a gazillion and one, a gazillion and two . . .,” and she describes the spontaneous invention of arguments that there is no highest number as common. In my exposition, I assume the subject has acquired a numeral system for counting in order to avoid complications with number-languages that are outstripped—at least in this respect—by the subject's knowledge of number itself. In any case, children in our culture do, in time, acquire the understanding of numeral systems. That achievement, too, is arduous because of confusions that can arise from the child needing to identify the terms of a numeral system with the number-words in natural languages; apart from this there is evidence that different brain circuits handle these different number systems.

¹⁰Kripke, 16: “Quounting” a heap of objects is counting in the ordinary sense unless the heap is formed as the union of two heaps, one of which has 57 or more items (assuming the subject has never so-far counted heaps that large), in which case the answer is 5.

¹¹So, for example, Carey (2009) offers an “ontogenetic” description of how the child—already possessing certain innate subpersonal cognitive systems (the parallel individuation of small sets, and the natural language quantifiers) that can be applied to particular (restricted) numerical tasks—is enabled over the course of a year and a half by induction

and analogy (“Quinean bootstrapping”) to understand certain crucial properties of numerals, e.g., that they continue indefinitely, and to understand how to apply them to count collections of objects. As we will see, a striking effect of the rule-following paradox is that Carey's empirical hypothesis, and the others like it that are offered by cognitive scientists, fail to even be *relevant* to a solution.

¹²Some of these possible “mistakes” actually occur during a child's acquisition of counting but most don't. Children go through specific stages as they acquire counting skills, and cognitive scientists describe children successively, as “one-knowers,” “two-knowers,” “three-knowers,” “four-knowers,” “subset-knowers,” and finally, “cardinal-principle knowers.” They are described—at one stage—as “one-knowers” because they know how to distinguish one object from many, but they can't distinguish between the cardinal numbers of groups of objects larger than one; “two-knowers,” are analogous, and so on. At one stage in their acquisition of counting skills, children also skip numerals when counting; this isn't something they do later. See Carey (2009), Butterworth (1999) for descriptions, and for references to the literature.

¹³“First I group them into collections of five because they're easier to recognize. Then I count all the groups, and multiply the result by five.” This is a description that a child might give because it strikes her as obvious that her new method gives the same answers that the ordinary way of counting (strict enumeration) does.

¹⁴As I'll indicate later, although this recognition-requirement seems to apply to counting and to other simple computational concepts, addition, subtraction, etc., it isn't a requirement on every concept that we take ourselves to understand. On the contrary: It's often thought that the meaning of a concept is understood by someone (or by all of us) even though he (or all of us) may not be able to recognize when some (or even all) of the uses made of the concept are wrong. See section 7 for discussion of this.

¹⁵Kripke, 7-22. Kripke poses the sceptic's challenge in first-person terms, as whether my previous practice and my occurrent thinking is compatible with “quaddition” instead of (as presumed) addition. I've instead described the challenge “third-person,” and—as noted—in terms of counting (instead of “quounting”). Given that we routinely allow ourselves to describe the phenomenology of third parties—pretty much as I've done above—this raises no complications. The reader uncomfortable with my use of our ordinary practice of describing the phenomenology experienced by others can easily restructure my discussion in first-person terms, and I invite him to do so. It won't affect the trajectories of my arguments.

¹⁶This is widely perceived to be the lesson of Wittgenstein's §139, and nearby sections. Putnam (1981), 20, puts the point very nicely: “What the phenomenologists fail to see is that what they are describing is the inner expression of thought, but that the *understanding* of that expression—one's understanding of one's own thoughts—is not an *occurrence* but an *ability*.” This isn't always noted, but even if a subject faces what strikes us as exactly the same counting task a second time, it's possible for her to do something different and yet for her to describe it as her doing the “same” thing she did before. (After all, the second task is occurring on a Tuesday instead of on a Wednesday, or during a full moon instead of during a new moon, or it's just occurring *later in time*.) Kripke (52, endnote 34) implicitly acknowledges the point when he quotes Wittgenstein: “If I know it in advance, what use is this knowledge to me later on? I mean: how do I know what to do with this earlier knowledge when the step is actually taken?” (1956, I,§3, italics his).

¹⁷This is called "subutilizing" by cognitive scientists, and it seems to be restricted to about four items for adults and to three items for children. See Mandler & Shebo (1982).

¹⁸DIS 2 illegitimately and rigidly singles out a *single* method. As noted in section 2, a subject may have acquired other appropriate short-cut methods of counting. For purposes of discussion, I'll set aside attempts to refine DIS 2 to handle this because such refinements won't affect the overall dialectical trajectory of this paper.

¹⁹The *actual* details of our biology—e.g., brain organization, physical laws, and so on—are pertinent to the truth of the counterfactuals involved.

²⁰Kripke (17) writes: "Normally, when we consider a mathematical rule such as addition, we think of ourselves as guided in our application of it to each new instance." Also (Kripke, 10), "I follow directions . . ."

²¹This is an extremely important aspect of how we understand rule-following, and it's due to relatively deep facts about our (shared) psychology—in particular to the peculiar ways that we take account of (and fail to take account of) the subpersonal abilities that are relevant to our execution of tasks that we consciously undertake. This interesting interplay between the personal and the subpersonal will have no role in my initial models of private rule-following (my first four *Robinson Crusoes*); but it will be part of the realistic psychological aspects of the last Crusoe to be described (in section 12).

²²A routinely described psychological result is the insecure reaction of a subject in a situation—e.g., in a classroom—where very simple instructions have been given ("raise your hand when the teacher holds up a green card, and lower it when she holds up a red card"), when, after a few episodes of this, everyone else in the class (as previously arranged) conspiratorially violates the instructions.

²³I draw the terminology, "straight" and "sceptical" solutions from Kripke, 66-67.

²⁴I take myself (in this section) to be expounding Kripke's (74-93) discussion of assertability conditions. Douglas Patterson (9/24/09—email) has raised the concern that a general replacement of truth conditions (for contents or statements) by assertability conditions leads to the view that it's *sentences* "otherwise meaningless, that have assertability conditions." And so, "what I'm entitled in the community to do is to make a noise, say." He further suggests that such a view isn't coherent unless an illicit appeal to truth conditions occurs in the description of what people are licensed to do. I think—despite my eventual denial that truth conditions need be replaced by assertability conditions—that this is wrong. I agree that sentences (without necessary and sufficient conditions of application) are what—on this view—people are entitled to apply or refuse to apply. Two views are possible: (i) The pattern of entitlements of such sentences is their "meaning"; (ii) Because of certain conditions that "meanings" must meet (e.g., that they must result in necessary and sufficient conditions of application) and that such a pattern of entitlements doesn't meet, these statements don't have "meanings." In neither case, it seems to me, does it follow that an assertion practice (built on "entitlements") is incoherent, or that it needs supplementation with truth conditions. The replacement of truth conditions by assertability conditions—Kripke (86) describes Wittgenstein as claiming—doesn't rule out ordinary uses of "true" and "false." For exactly these reasons, as I indicate in section 13, one can still utilize—instead of assertability conditions—truth conditions, provided one understands them in a "deflated" way, rather than as requiring correspondence relations to facts. I'm suggesting, in other words, that Kripke (and perhaps Wittgenstein) has built into the notion of "truth conditions" that they are required to provide correspondence relations to facts (Kripke, 72: "A declarative sentence gets its

meaning by virtue of its *truth conditions*, by virtue of its correspondence to facts that must obtain if it is true" (italics are Kripke's)). I deny that "truth conditions" need to be so understood.

²⁵My thanks to Douglas Patterson (9/24/09 email) for noting some problems in my earlier discussion of these cases.

²⁶I'm assuming that Robinson Crusoe can invent words, and can even invent words for complex notions that he acquires—such as cardinalities of collections of objects—on the combined basis of his (innate) dispositions and his experiences. There may be reasons to doubt that humans can do this on their own, even if some sort of rich innateness hypothesis about human dispositions is nevertheless true. Regardless, the acceptable empirical assumptions (about this) aren't in dispute now. If Kripke's way of deriving objections to the private model of language from the rule-following paradox is right—I am focusing especially on his second and third requirements for a solution to the rule-following paradox—it doesn't matter how sophisticated we allow Crusoe's dispositions to be: the paradox arises in any case.

²⁷Actually, I've been describing Crusoe and his world from the perspective of our *own* language. Phrases like, "judging Crusoe's terms on their adequacy," involve evaluations that are to be made in *our* terms. What's involved in doing this is something that's explored further in later sections, especially in section 8 and section 15. Let's meanwhile treat this as a preliminary way of speaking of the isolated Crusoe and the issues he raises—a preliminary way of speaking that will later be reevaluated and perhaps corrected.

²⁸Because his numbers are finite, I won't claim he has *our* numerical concepts. On the other hand, it's not obvious I should deny this either—at least of the numerical concepts he evidently *does* have: e.g., the ones of collections he has the dispositions to count, and to order in magnitude. See section 13 for further discussion of this issue.

²⁹Wittgenstein is widely taken to have challenged such concepts in various respects. But the considerations he raises seem to stem directly from rule-following considerations, and so it would be question-begging to press them in this context. I table for future work other ways that philosophers have thought to undermine this apparent datum about (first person) uses of "pain" and similar words. I rely on this property of such words in section 9, and say a little more about how I think they should be understood to operate.

³⁰Complications arise if we try to sharpen this distinction because that requires distinguishing Crusoe's dispositions to apply a word from other dispositions he may have that conflict with doing so. It doesn't matter for the form my argument eventually takes whether this distinction is tenable or not, and so I forego any further discussion of it.

³¹Kripke (on behalf of Wittgenstein) denies this is a possibility for a Robinson Crusoe. He (112, footnote 88) writes: ". . . in the absence of Wittgenstein's sceptical paradox, it would appear that an individual remembers his own 'intentions' and can use one memory of these intentions to correct another mistaken memory. In the presence of the paradox, any such 'naïve' ideas are meaningless. Ultimately, an individual may simply have conflicting brute inclinations, while the upshot of the matter depends on his will alone." What Kripke goes on to write is valuable to take note of: "The situation is not analogous to the case of the community, where distinct individuals have distinct and independent wills, and where, when an individual is accepted into the community, others judge that they can rely on his response . . . No corresponding relation between an individual and himself has the same utility." I'll not be drawing the same conclusions as Kripke has offered here on behalf of Wittgenstein.

³²It might be thought, because “coconut” has the same extension as “coconut,” that he does have access to our word. But, if so, what would this mean to *him*?

³³Think of a map. Some maps, we think we can say, “perfectly” describe the terrains they characterize. But how does this go beyond: using the map can’t mislead us in some way? We think it does go beyond this because we can imagine, after all, a piece of paper with lines on it, and we can also imagine a terrain the contours of which the lines in the map fit perfectly. (The lines on the map are isomorphic to the contours of the terrain *modulo* the respective sizes of the map and terrain.) This *imagery*, however, faces exactly the same problem. Our “image” of the terrain is still a description of the terrain “without language or concepts,” one that we have abstracted away from the input of our (visual) dispositions; and we are taking ourselves to recognize that the map fits this supposedly wordless (and nonconceptual) characterization of the terrain. Considerations like these suggest that Crusoe 2’s problem is our problem. Many philosophers over the years have claimed—against the correspondence theory of truth, for example—that one’s words can’t be compared to reality. It’s not always obvious what such philosophers mean when they say this. One thing some of them might have meant is this.

³⁴Maybe God can’t (or shouldn’t) say this. After all, so saying requires (i) a translation of Crusoe 2’s numerical terms to God’s, and (ii) the treating of deviations of Crusoe 2’s applications of his terms from God’s according to the standards of God’s usage. The point isn’t that God can’t treat his usages as superior to Crusoe 2’s; it’s that perhaps there is no real fact of the matter about how Crusoe 2’s terms should be translated to God’s. (See section 13.) In any case, God can still say: Crusoe 2’s words aren’t picking out real differences and similarities among objects—differences and similarities that bear on his (Crusoe 2’s) well-being, and that’s why he’ll do badly. Similar points can be made about Friday’s decision to translate Crusoe 2’s terms to his, and fault Crusoe 2’s usages using his (Friday’s) standards, except—as I note momentarily—that Friday isn’t in the position God is in vis-à-vis Crusoe 2’s failure to detect the “real” differences and similarities among objects.

³⁵It might have been thought that the opening description of Crusoe 1 and his world in section 7 was a description of him from a “God’s eye point of view.” Perhaps it’s better not to attribute such powers to *us* any more than to Friday. Perhaps it’s better to say that we described Crusoe 1 from a “Friday’s point of view.” Then our talk of Crusoe 1’s dispositions perfectly fitting his environment translates to this: If we were on the island with Crusoe 1, we wouldn’t be able to make this Crusoe 1 into a money pump: we wouldn’t be able to exploit him.

³⁶My thanks to Mitch Green (oral communication—August 8, 2009) for raising this issue. He isn’t responsible, however (and I apologize for), any distortions in his concerns due to how I’ve mistakenly recollected them.

³⁷“Partially,” because there are complications due to our sometimes allowing our evaluation of someone’s success to include an evaluation of their values. Thus, we may deem someone a failure because (although given their values, they have succeeded) we regard the successful satisfaction of *those* values to be nevertheless a failure. These complications don’t directly bear on the role of “success” as it’s used in this paper.

³⁸I don’t mean to suggest, by my use of the word “unhealthy,” that such Crusoes have values that we have a right to ignore or rule out. Nor do I want to claim that such Crusoes couldn’t engage in cogent private rule-following despite their “unhealthy” values. Many such clearly could. Rather, I’m pursuing straightforward cases of Crusoes (with, that is,

fairly straightforward dispositions) only to keep certain complications in exposition to a minimum.

³⁹I owe my attention to this concern to Stephen Schiffer (oral communication—August 8, 2009). He isn’t responsible, however (and I apologize for), any distortions in his concerns due to how I’ve mistakenly recollected them.

⁴⁰I distinguish “feeling hungry” from “hunger.” We ordinarily describe someone as hungry when they need food, and feel that way. We don’t think people are “really” hungry when they feel hungry but don’t need food (e.g., when tired, or because of certain physiological conditions).

⁴¹Some may think that I’m helping myself to an apparatus—introspective concepts—that are themselves open to challenge along the lines of (one or another version of) Wittgenstein’s “private language argument.” I don’t agree; see endnote 28.

⁴²It seems that Crusoe 3, by somehow using any of his own words or his own concepts to refer to something beyond what he’s disposed to apply them to is attempting to “raise himself up into the air by pulling up on his own bootstraps.”

⁴³Lewis (1983). As just noted, many philosophers have followed in his footsteps. See, e.g., Sider (2009), especially 400-401, where suggestive language like “reference magnets,” and “ideal interpreters” are invoked to rhetorically help the suggestion along. See my 2000, especially Part III, §5, for my earlier arguments against this approach to reference.

⁴⁴Well, maybe it is. But I’m not going to press this objection now.

⁴⁵The world can be as neatly divided into natural kinds as one likes. Regardless, its divisions aren’t *semantic* facts about the world. There are no *semantic* facts about the world. A slogan: *The world is semantically inert*.

⁴⁶In my original (2000, Part IV, §2) discussion of this issue, I labeled this the “projection problem.”

⁴⁷It isn’t hot air to demand that an interpreter charitably interpret a speaker’s words according to that interpreter’s *own* lights. But that’s an entirely different matter. (I’m not endorsing such a constraint on interpretation, by the way: I’m pointing out *differences*.)

⁴⁸Ways of so evaluating alternative private-language practices aren’t needed by Crusoe 1, of course, because his experiences of success are ones he can’t improve on.

⁴⁹With, however, a caveat I describe in the next paragraph.

⁵⁰No doubt it differs, for example, for the different senses, and it involves various sorts of automatic subpersonal faculties.

⁵¹This is a grammatically correct question of English, just as “Does ‘*Chaque champignon est vénéneux*’ correspond to a fact?” is a grammatically correct question of English. My thanks to Douglas Patterson (9/24/09 email) for his complaints about an earlier formulation.

⁵²I don’t think this can be made to work, but I won’t pause now to show this.

⁵³And I’m among them. See my 2006, my *forthcoming*(a), and my *forthcoming*(b).

⁵⁴Some philosophers would—in denying truth to be a correspondence notion—also deny that *any* truths correspond to facts. But that isn’t required to enable the sentences of Crusoe 5’s language to have truth conditions. All that’s required is that not *all* his truths need to correspond to facts.

⁵⁵I describe the view that we so self-ascribe ignorance to ourselves *the broad ignorance claim*. See Azzouni (*forthcoming*(b)), chapter 2.

⁵⁶This is just the rule-following paradox again. The grounding facts are missing, and

God can't legitimately stipulate answers in lieu of those facts.

⁵⁷In this case the c-quotes are not signaling quotations of Crusoe 5's sentences but instead characterizations of his thoughts.

⁵⁸Recall endnote 30.

⁵⁹Recall endnote 19.

⁶⁰See Azzouni 2006, my forthcoming(a), and my forthcoming(b). I should note that this third point is—in a way—anticipated by Kripke's Wittgenstein. Kripke (86) attributes to Wittgenstein the view that talk of truth and falsity, and even the use of a "calculus of truth functions" is compatible with his sceptical solution to the rule-following paradox. That means, I must add, that this talk is compatible with a notion of truth *sans* correspondence or fact. The additional insight I attribute to the truth-deflationist is that a notion of truth *sans* correspondence or fact is actually all the "truth" anyone needs; that includes the need to use it in truth-conditional analyses of languages. Such a fully functional notion of truth is only mistakenly seen as requiring either correspondence or facts. The references to my work mentioned at the beginning of this footnote, I note again, are pertinent to making this case.

⁶¹This is, of course, the main reason Kripke's Wittgensteinian sceptical-solution to the rule-following paradox is so dramatically shocking.

⁶²Kripke describes his own misgivings about Wittgenstein's sceptical solution this way: "But may the individual doubt whether the community may not in fact *always* be wrong, even though it never corrects its error? It is hard to formulate such a doubt within Wittgenstein's framework, since it looks like a question whether, as a matter of 'fact', we might always be wrong; and there is no such fact" (Kripke, 146, italics his). Kripke then notes that he has avoided a more extensive discussion because in doing so he "might have to abandon the role of advocate and expositor in favor of that of critic."

⁶³Recall the passage in Kripke, 21-22, which begins, "Sometimes when I have contemplated the situation, I have had something of an eerie feeling . . ."

⁶⁴I owe much of the particularities of how I pose this worry to Douglas Patterson—9/24/09 email.

⁶⁵As always, the ways we first-person apply certain words—"pain," etc.—seem to be exempt from this.

⁶⁶We are aware, of course, that words can change in what they refer to (they can change their "meanings"). But we see this as a sporadic and somewhat slow (e.g., generational) process—the kind of process, for example, that William Safire notoriously denounced.

⁶⁷I'm here abstracting away from the rich complexity of natural languages, and—of course—from the more complex semantic theories that are available to handle that complexity. In doing so, we might, for example, restrict attention to simple statements of the cardinality of collections of objects.

⁶⁸Notice that "truth conditions," so understood don't correspond to how the speakers use their words "true" and "false." They do *not* treat their assertions at a time as true by virtue of their dispositions to so assert them. This is because—on the contrastivist view—they are presupposing a false picture of their language(s).

⁶⁹One thing we cannot do is characterize the truth conditions of statements in terms of dispositions of people to assert or deny statements when a "mistake" isn't being made. This would only be to again face the rule-following paradox at the level of semantic theory: the notion of a "mistake" can only have content in relation to the dispositions of

speakers at some other time.

⁷⁰A being, perhaps, whose experience of coconuts—epistemically speaking—is like our experience of *pain*.

⁷¹So it's probably wise to *drop* descriptions of words "carving the world at its joints," or words "fitting the world as it is" because they are treacherously misleading for metaphysical thinking. The vanilla "suitably applying to the world" glossed—as I suggest—in terms of maximally positive success curves is better.

⁷²If someone—counter to my suggestion in endnote 70—insists on continuing to use the idiom "fits the world," *this* is how it must be understood.

⁷³What if there are no maximally positive success curves (what if every success curve can be improved)? Here are two possibilities. First, that all the successively more positive success curves induce application-conditions that agree on some of the applications of some of our words. Then, defeasibly, those words refer to those agreed-upon items. The second possibility is that, for each item that a given application-condition applies a word to, there is a later (more positive) success curve, after which the application-conditions induced by later, and superior, success curves exclude it. Then the still-valuable talk of *reference* is underwritten by *no* correspondence facts.

References

- Azzouni, Jody. 1994. *Metaphysical myths, mathematical practice: The epistemology and ontology of the exact sciences*. Cambridge: Cambridge University Press.
- . 2000. *Knowledge and reference in empirical science*. London: Routledge.
- . 2006. *Tracking reason: Proof, consequence and truth*. Oxford: Oxford University Press.
- . forthcoming(a). 'Deflationist truth'. In Michael Glanzberg (ed.) 'Handbook on Truth', London: Blackwell.
- . forthcoming(b). *Talking about nothing: Numbers, hallucinations and fictions*. Oxford: Oxford University Press.
- Butterworth, Brian. 1999. *The mathematical brain*. London: Papermac.
- Carey, Susan. 2009. 'Where our number concepts come from'. *Journal of Philosophy* 106, no. 4: 220–254.
- Gelman, R. & Galistel, C. R. 1986. *The child's understanding of number*. Cambridge: Harvard University Press. (originally published in 1978).
- Goldfarb, Warren. 1992. 'Wittgenstein on understanding'. *Midwest studies in philosophy* 17: 109–122.
- Kripke, Saul. 1982. *Wittgenstein on rules and private language*. Cambridge, Mass.: Harvard University Press.
- Lewis, David. 1983. 'New work for a theory of universals'. In 'Papers in Metaphysics and Epistemology (1999)', 56–77. Cambridge: Cambridge University Press.
- Mandler, G. & Shebo, B. J. 1982. 'Subilizing: An analysis of its component processes'. *Journal of Experimental Psychology: General* 11–22.
- McDowell, John. 1984. 'Wittgenstein on following a rule'. *Synthese* 58: 325–63.
- Putnam, Hilary. 1981. *Reason, truth and history*. Cambridge: Cambridge University Press.

- Rosch, E. & Mervis, C. B. 1975. 'Family resemblances: Studies in the internal structure of categories'. *Cognitive Psychology* 7: 573–605.
- Rosch, E., Gray, W. D., Johnson, D. M. & Boyes-Braem, P. 1976. 'Basic objects in natural categories'. *Cognitive Psychology* 8: 382–439.
- Sider, Theodore. 2009. 'Ontological realism'. In Chalmers, Manley & Wasserman (eds.) 'Metametaphysics: New essays on the foundations of ontology', 384–423. Oxford: Oxford University Press.
- Wittgenstein, Ludwig. 1956. *Remarks on the foundations of mathematics*. Oxford: Basil Blackwell.
- . 1958. *Philosophical investigations*. Third ed. New York: The Macmillan Company. (trans. G.E.M. Anscombe).