

Chapter Four: Linguistic Norms and Future Behavior

1. Introduction

The purpose of the previous chapter was to show that the ascriptional practices that Burge appeals to can and should be incorporated into an equilibrium for our semantic terms. The purpose of this chapter is to argue that, once one accepts the admissibility of the ascriptional practices that Burge appeals to, there is little reason not to accept the ‘retrospective’ attributions of content that we commonly make. Once our general beliefs about meaning are brought into line with the ascriptional practices appealed to by Kripke, Putnam and Burge, there isn’t much left to prevent us from incorporating the ascriptional practices brought to our attention by Wilson. Most of the general beliefs that conflict with such ascriptions will have been given up or modified already. Indeed, it will be argued in section five that there is a sense in which the practices appealed to by Putnam just *are* the ones appealed to by Wilson, with the true nature of those practices only being made clear with the Wilson cases.

Incorporating such future directed ascriptional practices into any equilibrium for our semantic terms involves accepting a type of “temporal externalism” (hereafter “T-externalism”) according to which the future behavior of an individual or his society can affect the conditions for the correct application of his terms. This chapter will discuss, motivate, and elaborate various aspects of a T-externalistic position.

Before beginning, however, I should make two terminological points. First of all, when discussing T-externalism, talk about subsequent behavior “helping determine the term’s final extension,” or future use “making our concepts more determinate,” is hard to avoid. Nevertheless, it invites a misunderstanding of the proposed view. Namely, it suggests that our concepts change over time (that is, they go from being indeterminate to being determinate), rather than having the same determinate content attributed at all times. According to the T-externalist, a term’s meaning is determined by the totality of facts relating to its use, and such facts are not limited to what is epistemically accessible to us in

the present. Both past *and* future usage is relevant to the term's interpretation. Future behavior helps determine a term's extension, but it does so by constituting some of the facts relevant to its proper interpretation. Before the future decisions were made, the extension was not indeterminate. Rather, the extension was just 'unsettled,' since all of the relevant facts were not realized, or epistemically accessible. Future use can "make a term's extension more determinate" in an epistemic sense, just as adding a piece to a puzzle can make what the puzzle represents 'more determinate' without in any way changing what the puzzle represents.

Secondly, it should be noted that, in light of the views presented in this chapter, "extension" could be understood as ambiguous between (1) what the term refers to, and (2) how the term's use is 'extended' into the future. Since it is argued that (2) helps determine (1), context alone may not necessarily make clear which of the two uses is intended. Because of this, I will try to use the term exclusively in the first of the two senses given.

2 T-externalism: the ascriptional practice

The fact that future usage can affect the way we at least *attribute* thought and linguistic content has been discussed in chapter one. The gist of these cases is that our everyday ascriptions of content reflect a sensitivity not only to the agent's internal make up, and the structure of his physical and social environment (including the past history of use of the terms he is using), but also to 'accidental' developments in his terms' usage in the period subsequent to the moment of utterance. The cases in question will be recapitulated and expanded immediately below.

2.1 Wilson's Druids

For instance, imagine a speaker, Edwin, who is a member of an isolated community inhabiting an island on which the class of birds and the class of flying things are coextensive. He has a term "ave" whose putative extension includes these locally co-extensive classes, and he has beliefs such as "all and only aves can fly," and "all aves are

living things.”²⁷⁶ Both of these beliefs are true of the “aves” on the island, but it is clear that one must be given up if, say, a plane lands on the island and he encounters a flying thing that is not a bird. We can imagine that Edwin’s total belief set will be split in just this way, with half of Edwin’s ave-beliefs favoring one possible equilibrium, and half of his beliefs favoring the other. There may, however, be no way to tell in advance whether an equilibrium for Edwin’s use of “ave” will include planes in its extension or not, because the mode in which Edwin first encounters planes may determine the comparative entrenchment of his beliefs and attributions.

While a term’s putative extensions include everything one is disposed to apply it to, items to which the term has actually been applied tend to be more deeply entrenched.²⁷⁷ For instance, if Edwin first sees planes flying high in the sky, his belief that all flying things are aves will lead him to call the planes aves. Planes may thus already be entrenched within the putative extension of “ave” when Edwin first sees one land and realizes that it is not a living thing. As a result, he may, upon seeing the first grounded plane, give up the belief that all aves are living things, and thus reach an equilibrium in which “ave” means *flying thing*. On the other hand, if he first sees planes on the ground, his belief that all aves are living things will entrench planes within the putative *anti*-extension of “ave.” This entrenchment may cause him, on discovering that planes can fly, to reject the belief that all flying things are aves, and thus reach an equilibrium in which “ave” means *bird*.

In each case, Edwin will not see himself as having changed what he meant by “ave” in any way. He will describe himself as having discovered that he had always been wrong in thinking that all aves were living things, or as having discovered that he had always been wrong in thinking that all flying things were aves. He will not see his use as just developing along one of two paths that were available for him. Instead, he will see his

²⁷⁶ We can assume that the rest of Edwin’s community is in a similar position. Indeed, the example here is adapted from Wilson 1982, and Edwin is a member of Wilson’s community of “Druids.”

²⁷⁷ One can thus reach a middle position between understanding the content of a speaker’s utterances purely in terms of the *actual* causes of his behavior (as in Davidson’s position), and one that seeks to understand content in terms of the *subjunctive* causes of that behavior (as in Fodor’s position).

present use as being completely determined by his previous use. The order of appearance of novel phenomena may affect the comparative entrenchment of elements within the putative extension and general characterization, and thus determine which of a number of possible equilibria is reached.²⁷⁸

2.2 Grant's Zebra

Furthermore, how our present utterances are interpreted may be affected not only by our *own* future behavior, but also by the future behavior of *others* (behavior that may even occur after our death). To return to an example from chapter one, around 1820 Grant introduced the term “Grant's Zebra” for a type of zebra native to Kenya.²⁷⁹ A few years later, Chapman introduced the term “Chapman's Zebra” for a morphologically distinct type of zebra found in present-day Zimbabwe. Later still it was discovered that the two types of zebra interbred near the Zambezi river and that, morphologically, one gradually faded into the other. Grant's and Chapman's zebras both turned out to be a races of the species *Equus burchilli* (one race of which, the Quagga, is arguably not a type of zebra at all).²⁸⁰ This story seems representative of the way we typically describe such cases. However, if it is accurate, then (as with Edwin's use of “ave”) it is merely a *historical accident* that the term “Grant's Zebra” had the extension it did. If the taxonomists had investigated the area around the Zambezi river *before* they hit deepest Zimbabwe, then they probably would have ‘discovered’ that Grant's Zebra could be found through most of East Africa, gradually

²⁷⁸ In this respect, linguistic development is much like case law, which, notoriously, has this property. The principles according to which judges develop a set of precedents are often determined by the order in which new cases are brought to them. For a discussion of this aspect of case law, see Dworkin 1977.

²⁷⁹ Of course, we do not know precisely what general beliefs Grant had at when he introduced the term, but we will assume here that his beliefs were equally compatible with either set of developments within the linguistic practice he initiated. In particular, even if he had some tendency to understand ‘type of zebra’ as ‘species of zebra,’ the fact that there are not, technically, any ‘species level’ classifications which include only zebras gives some incentive not to understand ‘type of zebra’ as ‘species of zebra.’

²⁸⁰ For a discussion of how the folk classification “zebra” does not seem to map on to any well defined biological kind see Gould 1983.

changing into a different subspecies as it drifted south. In such a case, “Grant’s Zebra” would have picked out the entire species, not just the race found in Kenya.

2.3 Reaction to the practice

Such cases²⁸¹ suggest that, when we interpret the past use of other speakers (and even ourselves), we help ourselves to subsequent specifications that were not determined at the time of utterance. Many, however, are strongly inclined to think that, if they involve a commitment to T-externalism, such ascriptions *must* be mistaken. There is a temptation to think that what Edwin means by “ave” today *can’t* depend upon whether the plane he sees tomorrow is in the sky or on the ground, and what Grant meant by “Grant’s Zebra” in 1820 *can’t* depend upon whether other naturalists explored Zimbabwe or the Zambezi in 1823. The conflict we see in the Wilson cases is one in which many, if not most, people are willing to say that the ascriptional practice is mistaken.²⁸² That is to say, they think that we should hold on to the belief that the content of our thoughts and utterances is independent of future linguistic development, even if it means giving up many of our ascriptions.

It should be stressed, however, that when just a single one of the scenarios is presented, the attribution of the later determinate content to the earlier stage seems perfectly natural. We are usually faced with just one such scenario when we actually go on to get a more determinate conception of what we are talking about. The intuition that makes this attribution so natural is, in some sense, tied to the presupposition that there will not be two alternate scenarios -- which is why the presentation of both scenarios undermines the intuition that is so strong in the single case. It is an important part of our linguistic practice

²⁸¹ Wilson discusses Grant’s zebra, along with a number of other such cases, and argues that they occur “in virtually every case of enlargement of our world view through scientific progress.” (Wilson, 1982, 572.)

²⁸² Including (possibly) Wilson himself. Wilson’s case is tricky because he claims that in cases such as Grant’s use of “Grant’s Zebra” or Edwin’s use of “ave” the reference of the terms in question changes but their meaning does not change. It is not, however, entirely clear that this is a consistent position, since it would allow that two successors of a given community could both mean the same thing by a particular word, even though they manifestly used it to pick out different classes of things.

that the people participating in it typically assume that how their terms are applied in novel cases will be determined by past use. What the Wilson cases suggest, however, is that this widespread assumption is often false. This is, of course, a uncomfortable realization if the practice described is one's own, and it might suggest that language is one of those practices the understanding of which undermines one's ability to take part in it. Furthermore, since cases where there are multiple equilibria available for our terms involve a breakdown of our presuppositions about language, any account of what our terms mean in such cases will seem unintuitive in at least some respects.²⁸³ As a result, it is not obvious that there will be any completely satisfying way to describe cases where the putative extensions and general characterizations of our terms have multiple equilibria available to them.

As long as one understands speakers in terms of belief and truth (rather than remaining at the level of general characterizations and putative extensions), disallowing any contribution from future behavior leaves the extension of such terms indeterminate. On such accounts, what, say, Edwin initially meant by "ave" would be indeterminate between *bird* and *flying thing*, and the reference of his term would have become more determinate (and so, in a clear sense, changed) when he first encountered planes. The ultimate acceptability of a T-externalistic equilibrium will depend, to a large extent, on the comparative merits of such indeterminacy-based accounts. Such accounts are discussed in the seventh section of this chapter. While there may ultimately be no fact of the matter as to whether T-externalistic or indeterminacy-based accounts of content are to be preferred, it is clear that the latter have been more fully explored. The following few sections will focus on T-externalism and argue that this much neglected position, while perhaps not compulsory, is a surprisingly attractive option, perhaps even more attractive than its indeterminacy-involving competitors.

²⁸³ We see a similar phenomenon with cases of splitting in discussions of personal identity: our conceptual apparatus works under the assumption that there will be no splits, and no application of this apparatus seems entirely satisfactory in the splitting cases. As Lewis puts it in another context, "Rival philosophical theories have their prices, which we seek to measure. But it's all too clear that for philosophers, at least, there ain't no such thing as a free lunch." (Lewis 1983a 353.)

3 T-externalism: arguments for the view

3.1 Specifying one's commitments, and the diachronic division of labor

As mentioned in chapter one, there is an argument form in which, given a set of criteria assumed adequate for determining what our words and thoughts are about, one constructs a thought experiment in which two individuals are identical with respect to the proposed set and yet can plausibly be said to refer to different things, thus showing that more must be added to the proposed set of criteria. Putnam is famous for arguing that our relations to the physical environment must be added, and Burge makes a similar point with respect to the social environment. The Wilson cases can be seen as suggesting that the future linguistic development should be included in the set of what contributes as well. Still, while it may seem as if the three factors represent a gradual expansion of the relevant factors, it should be noted that one can accept any of these additions independently of any of the others. In particular, one can accept that future events can determine what we mean, without accepting that our social context affects what we mean in a similar fashion. For instance, the ascriptions that tied what Edwin meant to his future usage could be accepted by someone who did not allow social usage to affect what an individual meant. The Grant ascriptions could not be accepted by such a person.

Indeed, the Edwin ascriptions might be the more compelling of the two for precisely this reason. When what we have said in the past underdetermines what we have committed ourselves to, it seems most plausible to say that at least *our own* subsequent utterances can be taken to specify the commitments involved in our earlier ones. One can be given credit or discredit based upon the correctness of one's assertions, and the assertions one goes on to make in the future can help determine the correctness conditions by which the entire set is evaluated. For instance, Edwin's initial use of "ave" had many aspects that favored interpreting him as meaning *flying thing* by the term, but just as many favoring interpreting him as meaning *bird* by it. As a result, Edwin could have found good reasons for claiming

that his term was properly used in either of the two ways. Given this background, we allow that Edwin always meant, say, *flying thing* by “ave” largely because he understood *himself* this way. If he had understood “ave” as picking out *birds*, we would have accepted his judgment on that as well. Unless our previous usage and dispositions clearly commit us to a certain equilibrium, we are entitled to specify just what we are committed to with our words.²⁸⁴ The fact that Edwin could not, if pressed, find anything in his previous use or dispositions that justifies his choosing one equilibrium over another does not prevent him from having the right to do so. Philosophical self-exegesis is a clear case where this consciously happens,²⁸⁵ but it often occurs without the speaker realizing that he is resolving a potential ambiguity. Even when there is more than one equilibrium available for our terms, certain contexts will make particular patterns of resolution seem so obvious that alternatives are not considered. Edwin’s case has precisely this property. While two equilibria are available to him, the way in which he initially encounters his first planes will make the move towards one of these equilibria seem obvious.

However, once we allow that individuals have the right to specify their commitments this way, the fact that one can have one’s reliance on others deeply entrenched in the general characterizations of one’s terms can result in the equilibria for one’s *own* terms being dependent upon the future behavior of *others*. For instance, if Edwin’s later use can determine what he *always* meant by “ave,” and there are people who rely upon Edwin for their use of “ave,” then Edwin’s future use can determine what *they* always referred to by “ave.” If his neighbor picks up the term from Edwin, and died before Edwin first encounters planes, Edwin’s usage after his friend’s death could determine what that friend referred to by “ave” before dying.²⁸⁶

²⁸⁴ For a fuller discussion of this, see the discussion of methodological individualism in chapter three.

²⁸⁵ This phenomenon is discussed in Grice 1957, 222-3. ²⁸⁶ What we mean by our terms, just like whether we are *eudaimon* or not, is largely dependent upon what we do during our lives, but it is ultimately independent of what others do around us nor of events that take place after our death. For a discussion of how future events affect our happiness, see Aristotle 1985, 1100-1101.

Of course, we rarely have a particular expert in mind when we rely on community use. Rather, we understand ourselves as taking part in the various linguistic practices associated with our words. As discussed in chapter three, this belief is often deeply entrenched, and so our words are often correctly taken to have the semantic properties that the surrounding practices associate with them.²⁸⁷ Such practices are not naturally understood as limited to thin temporal slices corresponding to the moment of utterance. Rather, the very idea of a linguistic practice involves some notion of temporal extension. Indeed, the admission of the Kripke phenomena and the contribution of past use already commits one to a notion of linguistic practices extending through time.

Just as an individual's usage does, a linguistic practice can have a number of equilibria available to it. The practice need not actually be viewed as *changing* unless it settles on an equilibrium that was not a member of the set originally accessible to it.²⁸⁸ A practice can evolve as its characterization of a term's meaning is made more determinate, but as long as it is only made more determinate, it can still be understood as the same practice. If the practice remains the same, so does the meaning of the term dependent upon it. When we make an utterance, we often commit ourselves to future refinements in communal usage because these refinements determine just what linguistic practice our usage is (has always been) a part of. Our conception of our linguistic practice determines just what our reliance on that practice commits us to.²⁸⁹

It was argued in chapter three that an individual speaker can be held to expert use in part because of his assumption that he is talking about just what the experts are. In much the same way, the experts in a given society can be held to their successors' use in virtue of seeing their predecessors, their successors and themselves (and their successors' seeing them) as part of a continuous process of discovering the extension of the same terms.²⁹⁰

²⁸⁷ For another discussion of this, see Evans 1982, 387.

²⁸⁸ This topic is discussed further in chapter one.

²⁸⁹ Once again, this sort of issue is discussed in greater detail in chapter three.

²⁹⁰ "Darwin's claim that diversity of species is due to a process of natural selection ... has undergone considerable amplification, qualification, and extension, but there are very few scientists today who

The synchronic division of labor is manifested in our confidence that, while we may not know exactly what falls under the extension of a given term, somebody else does. The diachronic division of labor is manifested in our confidence that, even if nobody knows exactly what falls under the extension of a given term, once somebody did, or someday somebody will. A scientist may introduce a term, knowing full well that its extension cannot be picked out determinately, leaving this task to future scientists (or perhaps his future self). The history of terms such as “AIDS,” “hepatitis c,” or “human growth factor” are typical examples of this process. By reading present specifications into the past, we license future interpreters to read future specifications into the present. Rejecting the correctness of such future ascriptions would commit one to the incorrectness of many of the ascriptions that one makes in the present, and since such future ascriptions are compatible with one’s present utterance, there is no reason to take on such a commitment.

That there is a diachronic division of linguistic labor involving past usage is relatively uncontroversial, but it is a point that is all too easily forgotten. If what we meant had to be determined entirely by present usage, much of the name-using practice that Kripke drew our attention to would be ruled out. Present applications of a term can be understood as mistaken at least partially because what a speaker now means by a word he is using is typically understood as being a function not only of how he is presently using the term, but also of how he (and others) have used the term in the past. The claim that what we mean by our terms is not solely a function of *present* usage should, then, be fairly uncontroversial. Some might, however, be less willing to allow the contribution of future usage than they would that of past usage. After all, allowing past usage to determine what we presently mean would seem to involve forwards causation, which is fine, while allowing future usage to determine what we presently mean would seem to involve backwards causation, which is far from fine. This is, of course, a serious worry, and

would deny that *Darwin’s* claim (and not just the more sophisticated selectionist positions which have since been developed) is true.” (Rouse 1987, 152.)

sections 4 and 6 will contain a discussion of how T-externalism involves neither a commitment to backwards causation nor a commitment to any substantive theory of time.

Of course, the mere fact that there is a diachronic division of labor extending into the future does not entail T-externalism. If facts about present use already settled what we meant (i.e. if there were only one accessible equilibrium), then the diachronic division of labor could be understood as merely the case of the future generations bringing to light commitments that were already settled in the past. It is only once one realizes that past usage does not settle how future usage should develop that one can see how the claim that the future clarifications determine what we mean leads to T-externalism.

3.2 Determinacy

As discussed in chapter three, one's belief that one's terms have determinate extensions favors the contribution of one's present social environment. Since admitting future usage can make what would otherwise be indeterminate contents determinate, this same belief supports the contribution of future usage as well. Neither Edwin nor Grant understand themselves as meaning something indeterminate by their terms. However, unless future use is allowed to help determine the extension of our terms, their extension will often be indeterminate. The fact that an individual's, or even an entire society's, usage can, up to any given time, fail to pick out a determinate extension for an individual's term gives one reason for trying to expand the pool of evidence relevant to what the individual picks out with his terms. Given that we actually make ascriptions that help themselves to future specifications, future usage is an obvious candidate to be included within the expanded pool.

3.3 Shared languages

Just as it is an advantage of non-individualistic accounts of content that they accommodate the idea that members of a linguistic community mean the same things by their words, it is an advantage of T-externalistic accounts of content that they can

understand present and past stages of a linguistic community as talking about the same things with the same words. The continuity that we are preserving is not only between past and future stages of our society, but between our past and future selves.

Indeed, the T-externalistic position, by allowing for more continuity of content, allows for more continuity between our past, present and future psychological states. If one is inclined to think of personal identity in terms of psychological continuity, then the T-externalistic position is much friendlier to the thought that we preserve our identity through time.²⁹¹ In much the same way, one's having a substantial conception of preserving our identity through time will also make one more sympathetic to T-externalism. If one understood persons as successions of temporal slices that were merely more-or-less similar to each other, then it might seem that the commitments one such slice could undertake at a moment could not be affected, or be affected by, the actions of the other slices. On the other hand, if one understands people as preserving their identity through time, what one commits oneself to at a moment can be affected by the actions that one makes at other moments. If the locus of commitments is itself temporally extended, then the claim that its assertional commitments need not be settled at the moment of assertion will seem more plausible.

4 T-externalism: an elaboration of the view

4.1 Meanings and histories

Since T-externalism gives future behavior a significant role in determining what we presently mean, it should not be surprising that some of the formal apparatus designed to account for explicitly future directed statements can also help explicate a T-externalistic

²⁹¹ Assuming, of course, that one understands psychological continuity in terms of the preservation of 'wide' contentful states. 'Narrow' psychological continuity would be preserved equally well on either view. Continuity of narrow psychological states is, however, a less appealing criterion of personal identity than the preservation of more fully contentful states.

semantics. In particular, the Edwin and Grant cases will here be characterized in terms of the following:²⁹²

1. A *moment* (“*m*”) is a spatially complete but instantaneous event: all of nature idealized to a zero temporal thickness.
2. The *causal ordering relation*, $m_1 \prec m_2$. (This relation is taken to be reflexive, transitive, and anti-symmetrical.)
3. A *history* (“*h*”) is a maximal chain of moments.
4. *Our World*. The set of all moments that are connected to this very moment by means of any zigzag combinations of the causal ordering relation or its converse.

The structure of Our World has the following features:

5. Historical connection: every two moments of Our World have some common historical ancestor.
6. No Backwards Branching: all branching is forward, never backward. Incompatible moments in Our World never have a common upper bound.

A bare moment is not enough for the evaluation of the truth of all of the utterances made at that moment. What one needs is a moment/history pair. The truth of many of our utterances, especially those about the future, depends upon how history continues after the moment of utterance, and so such histories must be included in the context of evaluation. The truth of a sentence is always relative to a history.²⁹³

Utterances that have their truth-value established at a moment, independently of what happens afterwards, can be described in terms of “settled truth.” Roughly, an assertion is *settled true* at a moment if it is true in *every* history passing through that moment.²⁹⁴ In much the same way, a term’s extension can be said to be ‘settled’ at a moment if the term has the same extension in every history passing through that moment.²⁹⁵ This is only a rough and ready characterization of what it is for an term to have a settled extension, since

²⁹² The formal framework used here draws heavily upon the Prior/Thomason tense logics (see Prior 1967, Thomason 1984), particularly as presented in Belnap and Green 1994. It should be stressed that the use of the Belnap and Green framework should not be taken to imply that I endorse their claim that the future is, in fact, open. Nor should it imply that they accept anything like T-externalism. Indeed, they seem to reject it explicitly (Belnap and Green 1994, 382).

²⁹³ See Belnap and Green 1994 373-4.

²⁹⁴ See Belnap and Green 1994, 374.

²⁹⁵ Or at least those futures where the meaning doesn’t change. And it may be an irreducibly ‘interpretational’ matter to decide which of these histories involve changing rather than simply developing the term’s meaning. This applies to our looking back on our own past history of use as well.

it assumes that, by the end of each history, each term's extension will have been settled. There will, however, be histories where a term's usage dies out before it is fully specified, and to account for such cases, one will need to modify the account of what it is for a term to have a settled extension as follows. A term will have a settled extension at a moment if either (1) there is at least one history containing that moment in which an equilibrium is reached and in every history containing that moment where an equilibrium is reached, the same equilibrium is reached (this should account for moments where a term has a settled extension even though there are some histories containing it in which no equilibrium is reached), or (2) the extension of the term in question has already been settled at some previous moment in all the histories containing that moment. The meanings of many of our terms may be 'settled' in just this fashion. In spite of this fact, one should not equate 'real extension' with this notion of 'settled extension' (which would suggest that if something is without a settled extension at m_i , then it doesn't 'really' have an extension at all). 'Real' extensions are always relative to a history.²⁹⁶

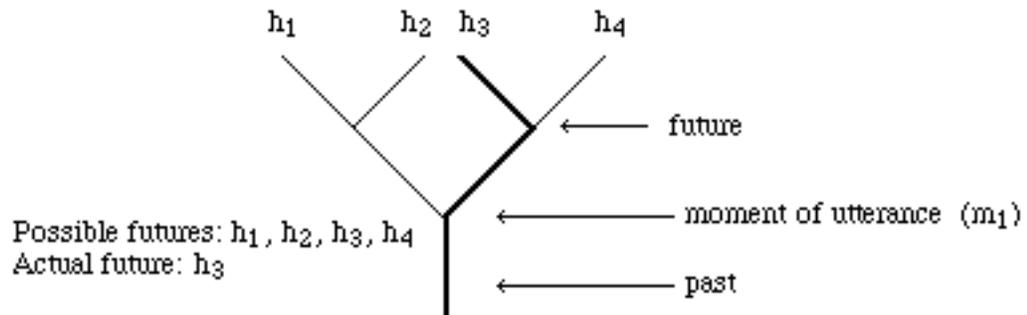
While the contribution of the future in determining the truth-value of statements either in the future tense or containing temporally loaded expressions such as "the losing bet" or "the final battle" has long been recognized, the T-externalist suggests that the class of these statements will be much larger than initially thought. In particular, the contribution of the future is often needed not only to determine the truth of the proposition expressed, but also to determine just what proposition *is* expressed. It should be noted, however, that even if the interpretation of a speaker's words may not be settled at a particular moment, the truth value of his utterance may be. Even if what Grant meant by "Grant's Zebra" depended upon how future usage developed, his utterance of "My zebra has stripes" will be settled true when uttered, since it is true on either candidate interpretation. This will not, however, always be the case. For instance, the zebras in Zimbabwe do not have striped feet, but

²⁹⁶ Belnap and Green (Belnap and Green 1994, 374) argue against equating 'real' with 'settled' truth, and if one sees a close relation between truth and meaning, one will expect a structural analogy here. If a term doesn't have an 'settled' extension at m_1 , it should not be assumed that it is then 'settled' that it has an indeterminate extension. ²⁹⁷ Lewis 1974, 110.

thus *always* settled that the term “Grant’s Zebra” would be applied to just the zebras in Kenya. A determinist thus need not see the Wilson cases as presenting a type of semantic indeterminacy at all. Such indeterminacy was supposed to arise because there were two alternate ways in which the usage could develop, but if determinism is true, there is no real alternative to our actual development. The future contains a set of settled facts that, while epistemically inaccessible to us at present, remain settled facts nevertheless. Allowing future events to contribute to the present meanings of our terms would thus not prevent us from taking there always to be settled facts about what we mean. It would thus seem less unintuitive to say that future events contribute to the present meanings of our terms. Some of the relevant facts may be epistemically inaccessible at the moment of utterance, but they remain settled facts nevertheless. A deterministic T-externalist can follow Lewis in insisting that we should not be interested in how *we* determine the facts, but rather in how *the facts* determine the facts.²⁹⁷

4.3 The thin red line

Much the same can be said if one accepts indeterminism with the thin red line (hereafter TRL). Indeterminism involves there being just one past, but a number of possible futures (represented here by a tree structure with a common trunk leading up to the present and various branches extending from it into the future). TRL assumes that, while there are a number of ways in which the future *could* possibly develop, there is still a fact of the matter as to what *will* happen. The possible future that is also the actual future can be understood as a “thin red line” (represented here as a thick black line) passing through the tree representing the possible histories.



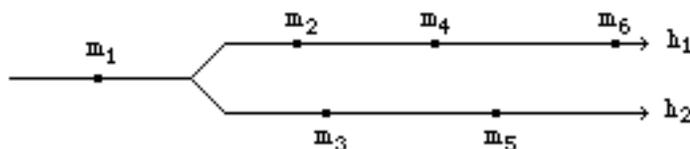
The T-externalist can describe the Grant's Zebra case with such a picture of the future as follows. Let us assume that Chapman discovers the zebras in Zimbabwe before explorers find them on the Zambezi on h_1 and h_3 , while the order of discovery is reversed on h_2 and h_4 . On such a picture, we can say that, at m_1 , while it remains *possible* for the explorers to investigate the zebras by the Zambezi first, they will, *in fact*, see ones in Zimbabwe first. With TRL, the admission of indeterminism does not prevent there from being a fact at m_1 about what will happen after it. The notion of epistemic inaccessibility may be stronger here, since with determinism it is at least in principle possible to discover how future usage will develop by looking at how things stand at m_1 . If determinism were true, knowing the states of all the sub-atomic particles at m_1 and the laws of physics would allow one to know, in principle, how the usage will develop. A Laplacean demon could always calculate future use based upon its knowledge of present facts. No such calculation could be made with TRL. Still, in many important respects, TRL is like determinism. While it is not a fact *about* m_1 that zoologists will not go on to apply "Grant's Zebra" to all of *Equus burchilli*, it is still a fact *at* m_1 that they will not do so. By contrast, on the deterministic picture, it is both a fact *about* m_1 that they will not go on to apply "Grant's Zebra" to Zimbabwean zebras, and a fact *at* m_1 that they will not do so.

With both determinism and TRL, the future events just reveal what always was to be. All the facts constitutive of the meaning were already in play, and thus so was the meaning itself. The passing of time simply makes open to us facts that were previously epistemically inaccessible. This makes the ascriptions associated with Wilson's thought

experiment more palatable. If it was already true at m_1 that Chapman would (in fact) discover the Zimbabwean zebras, before any zoologist hit the Zambezi, then it seems less strange to say that Grant meant just Grant's Zebra (and not *Equus burchilli*) by "Grant's Zebra" even at m_1 .

4.4 The open future

On the other hand, if the future is "open," then, at the time of Grant's initial dubbing, there will be not facts about how usage will develop, and thus about which equilibrium will be reached. A T-externalist who took the future to be open would thus have to describe the Grant's Zebra case as follows:



Treating m_1 as Grant's initial use of the term, there are two possible histories, h_1 and h_2 , in which the use of the term develops in different ways. In *our* history, h_1 , the term "Chapman's Zebra" was introduced at m_2 , the interbreeding discovered at m_4 , and we are discussing Grant's utterance at m_6 . In the alternative history, h_2 , the gradual shift at the Zambezi was discovered at m_3 while the Zimbabwean zebras were found at m_5 . On h_2 , "Grant's Zebra" denotes all of *Equus burchilli*, while on h_1 it denotes just Grant's Zebra.

T-externalism with the open future thus leaves it at m_1 not only *epistemically* but also *metaphysically* unsettled what Grant means. It is neither a fact *about* m_1 nor a fact *at* m_1 that the "Grant's Zebra" using practice will develop in any particular way. However, while some might find this fact a little 'unsettling,' it seems irrelevant to our *present* interpretation of Grant's 1820 utterances containing "Grant's Zebra." There are *now* settled facts about Grant's society's later history just as there are settled facts about its prior history. While it may be the case that, at the time of Grant's utterance, it wasn't *settled* that this would be the history that became the actual one, relative to *this* history, h_1 , Grant's term *always* meant

Grant's Zebra (I.e., it is now settled that he meant Grant's Zebra at m_1 , though it was not settled at m_1 that he meant Grant's Zebra.) At m_1 the full context of evaluation was not in place. Furthermore if there really is an open future, then evaluations of our utterances will have to be relativized to possible histories *anyway*. We would look back on Grant's future tensed utterance of "I will never be discussed in anyone's dissertation" and say that it was false, even though, at the time of utterance, it was not settled false. In an analogous fashion, we should look back at his utterances involving "Grant's Zebra" and treat them as referring to Grant's Zebra, even though, at the time of utterance, that extension was not yet settled.

5 T-externalism: three further consequences

There are a number of interesting consequences and incidental payoffs that result from the adoption of a T-externalistic position. Issues that seem puzzling within a temporally bound framework can lose much of their paradoxical character once future use is brought into play. Three such issues will be discussed briefly here. Once one allows future use to contribute towards what speakers meant in the past, the so-called "causal theory of reference" for natural kind terms can be stripped of some of its more dubious metaphysical and psychological presuppositions.²⁹⁸ We can say that we have always meant what we now do by terms such as "water" or "gold" without having to say that there is some special connection between linguistic usage and natural kinds. The initial baptism need involve neither unrealistic 'referential intentions' on the baptizer's part nor the possession of metaphysically privileged 'elite properties' on the part of the object baptized.²⁹⁹ Our linguistic precursors did not pick out gold by their term corresponding to "gold" because mineral kinds like gold are somehow 'more real' than functional kinds such as "yellow metal." Our precursors almost certainly had no explicit intention to pick out a natural (rather than, say, a functional) kind when they introduced the term. However, their vague

²⁹⁸ For a discussion of some of the problems with such presuppositions, see appendix A.

²⁹⁹ The first of these is appealed to in Putnam 1981d and Kripke 1972, the latter in Lewis 1983a.

presuppositions about the term probably made equilibria corresponding to both natural and functional kinds accessible. The compatibility of their usage with a number of equilibria allowed us to apply our own specifications retroactively. T-externalism allows one to incorporate an account of natural kind terms within a diachronic social picture, thus allowing one to provide a unified account of both the Putnam and the Burge phenomena.

T-externalism also lends itself to a fairly quick and cheap response to the skeptical paradox Kripke attributes to Wittgenstein. Kripke's skeptic assumes that if I meant something in the past, then it must be in virtue of some *past* fact about me.³⁰⁰ The skeptic must initially admit that we *now* mean *plus* and not *quus* ("otherwise, we will be unable to *formulate* our problem"),³⁰¹ and the T-externalist can say that *this* is precisely what makes it the case that we meant *plus* in the *past*. Of course, this present stipulation would do no good if it just amounted to my present utterance of the noise "in the past I meant plus," since such an utterance would, itself, be subject to interpretation. However, given that present use is not in question, one can help oneself to fully contentful use in the clarification of what one meant in the past. When Edwin initially grasped the rule for "ave," his intention and the rule he grasped did "determine a unique answer for indefinitely many new cases in the future."³⁰² However, it was only settled just what this rule was once he became acquainted with airplanes. Past intentions do 'determine' correct answers for new cases in the future, but the determination here is more 'logical' than causal. The content of past intention is determined by the answers we give in the novel cases.³⁰³ There can be, as it were, an 'internal connection' between how the speaker applies the term in novel cases and the rule he is taken to have grasped in the past.

³⁰⁰ Kripke 1982, 13.

³⁰¹ Kripke 1982, 14.

³⁰² Kripke 1982, 8.

³⁰³ Brandom (1989) argues that Hegel defends a similar view about the relation between our intentions and actions; that is, the content of the intentions we originally form is determined by the actions that those intentions result in. Just as becoming clearer about what one's practical commitments are as one acts need not be understood as a change of mind, clarifying one's theoretical commitments need not constitute a change of meaning as long as one stays within the range of accessible equilibria.

Still, while Kripke's arguments cannot show the T-externalist that I didn't mean *plus* by "plus" in the past, they could, if good, show that it wasn't *settled* that I did. One does not want to say that, if I had just recently given "5" as the solution to "68 + 57," then *quus* would have been what I *always* meant by "plus." The quick and cheap response alone would not prevent each new application of a term from being an "unjustified leap in the dark."³⁰⁴ The fact that such leaps in the dark 'retroactively' make it the case that they are in accord with the rule does not change the fact that, if Kripke were right, previous use would not 'guide' or 'constrain' us with respect to novel cases. While this can *sometimes* be true (as in Edwin's use of "ave"), it is not *always* the case. One wants, then, to be able to say that while Edwin's use of "ave" may leave it open whether the term applies to planes or not, there does not exist a similar choice as to whether the term applies to, say, robins (or any other species of flying bird not previously encountered on the island). Appealing to the need to reach an accessible equilibrium (of the sort discussed in chapter one) between his term's putative extensions and general characterizations will be enough to do just that.³⁰⁵

Finally, T-externalism points towards a substantive version of the Pragmatist view that truth is "the opinion which is fated to be ultimately agreed to by all who investigate."³⁰⁶ Peirce's maxim has often been criticized for lapsing into either falsity or emptiness once the notion of the opinion "fated to be agreed upon" is probed further. Nevertheless, a substantive, and still plausible reading of Peirce's maxim can be given in light of the considerations above. Namely, it can be understood as a conclusion about meaning, and not particularly about truth *per se*.³⁰⁷ The meanings of our terms become settled through the process of inquiry, then one could plausibly say that the 'end of inquiry' is the ideal limit where an equilibrium is reached for all the relevant terms. Until this point is reached, the truth values of many sentences will remain unsettled. Furthermore, since (especially

³⁰⁴ Kripke 1982, 10, 15.

³⁰⁵ For a discussion of how this could be done, see my "We Live Forwards but Understand Backwards" (MS 1995).

³⁰⁶ Peirce 1878, 273.

³⁰⁷ William James' forward looking account of intentionality is especially hospitable to T-externalism. (See my "James' Pragmatic Theory of Intentionality" MS 1995.)

with the open future) there is no normative constraint that completely determines what the end of inquiry will be, the claim made is a substantive one. On different histories, inquiry will reach different ends and claims such as “Grant’s Zebra has striped feet” will turn out to have different truth values. Assertional commitments are determined by the process of inquiry, and so truth-conditions depend upon the end of inquiry.

There are thus a number of advantages to adopting a T-externalistic framework. There may, however, be a number of disadvantages as well, and such potential problems are the topic of the next section.

6 T-externalism: some objections and replies

This section will consider, and attempt to answer, a number of objections to the suggestion that future usage can contribute to the present content of our thoughts and utterances.

Objection: T-externalism seems to endorse a type of backwards causation by allowing future events to determine the content of what was thought and said in the past.

Reply: While T-externalism allows future events to determine what we meant in the past, the type of determination involved is not causal. As a result, no backwards causation is involved. According to the T-externalist, the fact that the practice within which a given word is used extends through time allows future usage to help constitute the very practice that the speaker was taking part in at the moment of utterance. The speaker’s utterance is understood in terms of a characterization of the practice that may not have been available at the moment of utterance, but the characterization remains correct nevertheless. T-externalistic characterizations no more involve backwards causation than do claims like “John bought the winning lottery ticket on Tuesday.” John’s ticket may only have become “the winning ticket” when his number was drawn five days later, but one does not require

any notion of backwards causation to claim that he bought the winning ticket on Tuesday.³⁰⁸

Objection: If T-externalism does not understand the relation between present contents and future use as a causal one, this itself presents a problem for this view. Externalist theories such as Davidson's perceptual externalism exploit precisely this causal connection between thought contents and the 'external' factors that are meant to affect it. If future usage has no causal affect upon present utterances, then it should not effect the content of these utterances.

Reply: While T-externalism does not involve any backwards causation, it doesn't mean that we cannot understand there to be any causal relations between present utterances and future usage. There is still a causal relation between the two. Not only can present utterances causally affect future behavior, but the past use that affects our present utterances can be described as being part of a practice that also includes future usage. We can thus say that while future usage does not causally affect our present utterances, it can help determine the relevant characterization of the past usage which *does* causally affect our utterances.

Objection: Unlike other types of externalism, T-externalism seems vulnerable to a paradox similar to one associated with the idea that courts can fill in 'gaps' in the law *ex post facto*. For instance, imagine that, in 1980, a couple decides to get married while on a cruise ship. The ship's captain starts the wedding ceremony when they are still at sea, but by the time the ceremony is over, the ship is at port. The upon disembarking, the putative husband is killed by a passing car and, in a battle over his estate, his relatives contest the

³⁰⁸ Worries about backwards causation and T-externalism may bear some resemblance to worries arising from Aristotle's discussion of how we could be subject to "posthumous harm" (Aristotle 1985, 1101). In much the same way, the claim that we are subject to posthumous harm should not be understood in causal terms. Rather, it seems better understood in terms of future events making available descriptions of our life that make it seem less choiceworthy.

legality of his marriage. The law clearly states that ship's captains have the right to marry a couple at sea, and that they have no such right while at port, but it says nothing about whether they have such a right if the ceremony is begun at sea and ended at port. The case eventually reaches the state court of appeals in 1992, and the court decides that the marriage was, in fact, legitimate. If the courts count as filling in such gaps in the law *ex post facto*, then we can say that court *made it the case* that the couple was married in 1980. This seems to lead, however, to a paradox. If they were married in 1980, then, by the time the court made its decision in 1992, the couple had already been married for twelve years, and if they had already been married for twelve years, how could the court then "make it the case" that they were married. How can one *make* the case something that already *has been* the case for over ten years? One shouldn't be able to *make* something so if it already *was* so.

A similar paradox seems to arise for the T-externalist. Suppose that Edwin sees his first plane in 1990 at the age of thirty three, and that his use of "ave" between 1990 and 1991 makes it the case that he means *bird* by "ave." The T-externalist also seems committed to saying that his usage also makes it the case that he *always* meant *bird* by "ave." However, if Edwin *always* meant *bird* by "ave," then he meant *bird* by "ave" in 1970, and so by the time he had seen his first plane, he had meant *bird* by "ave" for over twenty years. How, then, can what Edwin did in 1990 make it the case that he meant *bird* by "ave" if he had already meant *bird* by "ave" for over twenty years?

Reply: Such paradoxes arise to a large extent because of the comparatively course-grained character of the phrase "make it the case that." It does seem paradoxical to claim that one made something the case at time *t* if it already *was* the case *before* time *t*. Such paradoxes disappear, however, if we describe Edwin's usage in 1990 as *settling it* that he meant *bird* by "ave," rather than *making it the case* that he did so. Talk of settling makes implicit reference to other possible histories, and so does not involve the same sort of commitment to what could not have been the case in the history in question. If Edwin

made it the case in 1990 that he meant *bird* by “ave,” then it could *not* have been the case that he meant *bird* by “ave” in 1989. However, if he only *settled* it that he meant *bird* by “ave” in 1990, it could still be the case that he meant *bird* by “ave” in 1989. It just couldn’t have been *settled* that he did. That is to say, there must have been possible histories passing through the moment in 1989 in which he did not mean *bird* by “ave.”

Describing what Edwin did in 1990 in terms of settling it that he meant *bird* thus does not produce the paradox associated with claiming that he made it the case that he meant *bird*. Nevertheless, isn’t the T-externalist still committed to saying that what Edwin did in 1990 made it the case that he meant *bird* by “ave” in 1960? It is not simply enough to find an alternative way of describing the cases which does not produce the paradox, the T-externalist must also show that he is not, in fact, committed to the style of description that does produce the paradox. Does the T-externalist have the resources to *deny* that what Edwin does in 1990 makes it the case that he always meant *bird* by “ave”?

To put himself in a position to make such a denial, the T-externalist needs to give an account of what *does* make it the case that Edwin means what he means, and then show how such an account prevents one from claiming that what Edwin did in 1990 made it the case that he always meant *bird* by “ave.” The T-externalist is, in fact, positioned to do just this. According to the T-externalist, what Edwin means at a given time by a term is determined by his entire temporally extended practice of using that term. As a result, what Edwin means by “ave” in both 1970 or 1990 is determined by the entirety of his usage from 1960 through 1992. What makes it the case that he means *bird* by “ave” from 1960 through 1992 is his entire history of using the term from 1960 through 1992. His 1990 usage is certainly part of that history, but it is not by itself something that makes it the case that he meant *bird* by “ave” in 1970. It *is* the case that he meant *bird* by “ave” in 1970, but what makes this the case is the fact that he meant *bird* by “ave” from 1960 through 1992. Stretches of usage make it the case that one means something by a term for stretches of

time, but moments of usage do not serve in explanations of what one means for moments of time.

Objection: What one means cannot extend beyond what one could communicate, so information that is unavailable to one's interpreters cannot contribute to what one means. Future use is unavailable to one's interpreters, so it can't contribute to meaning.³⁰⁹

Reply: First of all, one could argue that there is a sense of "communicate" in which one can say that what was *communicated* at a time was also determined by future usage. After all, if Edwin tells a friend that he killed an axe on his twenty first birthday, then that friend can come to believe that Edwin killed an axe on his twenty first birthday. According to the T-externalist, the content of both Edwin's utterance and his friend's belief will be determined by future usage, and so it is not much of a stretch to claim that what Edwin communicated by his utterance was so determined as well. The objection thus rests on a particularly restrictive notion of meaning as a communication by which both are limited to what could be determined by a contemporaneous interpreter.

However, there is little reason to restrict the evidence for an interpretation to what is available to a contemporaneous interpreter, since doing so would exclude important features of the society's *past* use as well. Such a restriction would thus rule out much of the name-using practice that Kripke describes along with the T-externalistic ascriptions. Since theories limiting the scope of what we could mean to the epistemic horizons of a contemporaneous interpreter are not nearly as well entrenched as the backwards-looking ascriptions they conflict with, it seems fairly clear that one should give up the restrictive theory. The inability of an interpreter to determine what a person means given all the facts

³⁰⁹ Something like this reason might be what leads Wilson not to draw T-externalistic conclusions from the druid and zebra scenarios. Wilson takes it to be a "rough adequacy condition" on any theory of meaning that "The evidence for assignment of an extension to a predicate should be limited to such linguistic behavior as can be reasonably extrapolated from the community's contemporaneous practice and should not reflect accidental features of the society's later history." (Wilson 1982, 553.) It is certainly true that, if we accept the first conjunct of Wilson's condition, we must accept the second. However, for reasons given in the reply, the first conjunct cannot be accepted.

about him at a given time should be taken as showing that all the “semantically relevant facts” need not be available at that time.³¹⁰

One might, of course, try to accommodate the Kripke ascriptions by allowing past and present, but not future, usage to contribute to what we mean.³¹¹ Nevertheless, one should have some principled reason for doing so, and once it is clear that causal influence isn't the issue here, the clearest motivation for this ‘asymmetry thesis’ disappears. Principled reasons of a quasi-internalist sort could be given for ruling out both, but it is much less clear how one is to justify ruling out one without the other. (Though reasons of a vaguely verificationist sort might be marshaled in favor of looking at the contribution of the future in preference to that of the past.) For instance, many of the intuitions that support the contributions of past usage (e.g. those of determinacy and continuity of meaning) support the contributions of future usage as well.

One might think that one could rule out the contribution of future usage because, if the future is open, facts about future use are not even *in principle* available to the interpreter. However, when we *now* interpret, say, Grant's original utterance, it is patently false that we don't have access to his terms' subsequent usage. What information is available to interpreters depends upon which temporal position one takes them to be interpreting from. The claim that what we mean must in principle be available to an interpreter has some plausibility, since what we mean by our terms is a quite plausible candidate for something which should not extend beyond what can be known. However, such arguments tying a range of facts to what can be discovered lose most of their plausibility if accompanied by such *ad hoc* restrictions upon the investigator's evidential base. As long as one allows an evidential base which includes future behavior, T-externalism is perfectly compatible with the requirement that what a speaker means by his terms be available to an interpreter. Finally, it should be noted that, future usage *is* in principle available to one's interpreter's,

³¹⁰ For a discussion of this, see Grandy 1973, 450.

³¹¹ Indeed, one might think that this was Wilson's real intent when he proposed his adequacy condition. However, Wilson's positive account of linguistic meaning does, in fact, rule out any contribution from past use.

they need only wait. On the other hand, barring the possibility of time travel, there may be no way for our interpreters to become acquainted with crucial aspects of past usage. As a result, future usage may be better off than past usage when we consider in principle availability to interpreters.

Objection: We don't know how future use will develop, so if such use contributes to what we presently mean, then we don't know the content of our own thoughts and utterances.

Reply: Given that we are not authoritative about the world's physical structure or the usage of our community, temporally-infected ascriptions no more threaten self-knowledge than the ascriptions Burge and Putnam appeal to. Self-knowledge creates no problems that are unique to T-externalism, and the strategies discussed in chapter three for reconciling self-knowledge with the Putnam/Burge phenomena will do so for T-externalistic phenomena as well.

Of course, the inaccessibility involved is of a different sort with future use, since one might think that there is, at least in principle, a possibility of the agent finding out what he means in the Putnam or the Burge cases. For the Burge cases the speaker could consult a dictionary or ask an expert, while for Putnam cases such as our use of "water" in 1700, one might think that the speaker could, in principle, perform the experiments and come up with the scientific theory needed to discover the true extension of "water." However, this difference is not one that makes a difference when what we are concerned with is self-knowledge. The relevant conception of self-knowledge is tied to notions like introspection, not to our ability to, say, go look up a definition in a dictionary. As a result, preserving self-knowledge gives us no new reason for rejecting T-externalism. Indeed, once again, given the practical inaccessibility of the past and the eventual arrival of the future, issues relating to self-knowledge should seem more pressing for those who allow the contribution of past usage than for those who allow the contribution of the future.

Objection: By allowing future use to determine the contents of our thoughts, T-externalism prevents us from making sense of the fact that our attitudes can be used to explain our behavior. All the facts having to do with whatever causally affects our behavior will have to be settled before that behavior occurs, so if T-externalism entails that the contents of our thoughts are not settled at the time of our action, then it seems to entail that our thoughts do not causally affect our behavior.

Reply: If T-externalism does spoil the role of attitudes in behavioral explanations, then they were spoiled already by the practices Burge and Putnam describe.³¹² As a result, if one accepts Burge's thesis about the social character of meaning, then problems with behavioral explanation will give one no reason to treat the practice Wilson describes as mistaken.

It should also be noted that T-externalism ascriptions can do a better job of 'making sense' of the speaker's behavior in the novel cases. We can explain why Edwin calls, say, the planes he initially sees in the sky "aves" by the fact that he means *flying thing* by "ave." By preserving the 'internal connection' between content and future applications, one can understand the speaker's behavior in novel situations as rule-governed in a way that one would not be able to otherwise.

Objection: By allowing future use to determine what we mean, T-externalism allows future developments that we would intuitively take to be *departures* from our use to determine what we 'really' meant all along. There seem to be clear cases where future development in a term's use should *not* be retroactively applied to the past. For instance, the term "girl" was once applied to both girls and boys, but one should not have to say that the fact that we now apply the term to just girls entails that previous applications of the term to boys were mistaken.

³¹² Once again, this is a subject discussed in greater detail in chapter three.

Reply: Before considering this objection, it should be remembered that an equilibrium is “accessible” if it does not require giving up elements that are too deeply entrenched.³¹³ Not all possible equilibria will be accessible. A term’s usage at a time may leave open a number of equilibria accessible to it, and these accessible equilibria set limits upon how much linguistic usage can change without there being a change of meaning. In order for it to be read back retroactively, future usage must develop within the range of equilibria that present usage has accessible to it. This severely constrains the types of linguistic development that can be read back into present use, and most perceived departures from present use will not satisfy such a constraint. If the changes in usage move to something that was not originally an accessible equilibrium, then there will be a change in meaning (as in the case of “boy”). On the other hand, if the change in usage can be understood as a move towards one of the accessible equilibria, there does not need to be any change in meaning (as in Edwin’s use of “ave”).

It is doubtful that any fixed procedure is available to determine when future use changes rather than clarifies what was meant before. Nevertheless, there should be a few rules of thumb about when a development is moving within the range of the term’s accessible equilibria and when it is not. Finding these rules of thumb should illustrate that our pre-reflective notion of ‘change of meaning’ does correspond to a recognizable class of cases within the framework sketched in chapter one.

Changes in a term’s conditions of application do not necessarily constitute a change of meaning. However, if enough of the original samples drop out of the putative extension (as in all the boys no longer being called “girls”), or if enough items from the term’s putative anti-extension (those items originally seen as definitely not falling under its extension) are added (beads not part of a rosary coming to be called “beads”) then, unless such radical shifts are justified by appealing to deeply entrenched aspects of the general characterizations, there may be a change of meaning.

³¹³ This is discussed in further detail in chapter one.

Radical changes in the general beliefs can occasionally cause changes in meaning as well. If we were suddenly to treat “water” as a functional kind term, we might be viewed as changing its meaning, even if the putative extension of the term remained the same (i.e.: if all the liquids we knew that had the required functional properties were made up of H₂O). However, it would constitute a change in meaning at least in part because the framework that was given up was a ‘viable’ one for “water.” (That is to say, there is at least one accessible equilibrium that incorporates the natural kind framework.) If investigation had proved that water was not predominantly a single kind, and we could no more find a natural kind corresponding to “water” than we could for “tree” or “lily” or “air,”³¹⁴ then we could reclassify it as a functional kind without changing its meaning.

Such an account would allow for considerable continuity of meaning through many radical shifts in our conception of the type of things we were talking about, but not all.³¹⁵ People who classified the animals in their environment in terms of a totemistic framework (members of certain tribes are descended from certain animals etc.) need not mean anything different than we do by their animal terms. This is because their most general, or ‘framework,’ assumptions no more determine what they mean by “goat” than our assumption that lilies formed a natural kind determined what we meant by “lily.” There is no accessible equilibrium incorporating the totemistic framework, so giving up that framework need not constitute a change of meaning for one’s terms. Things could be quite different with a people who classified things exclusively by their functional or aesthetic qualities (say, classifying birds by the type of meat they provided or the type of call they had). If they did not see the functional or aesthetic qualities as corresponding 1:1 with something like biological kind (or at least were willing to favor the functional classification if a conflict became apparent), then, even if they could be brought over to see the natural world in terms of ‘natural kinds,’ this conversion could be viewed as changing the

³¹⁴ For a discussion of how frequently our ‘folk-biological’ kinds fail to map on to any biological kind, see Dupré 1981.

³¹⁵ For another discussion of this, see the section on holism in chapter one.

meanings of all their terms, since the framework given up was a viable one. In any case, meaning change is not a problem unique to T-externalism, it is just the problem that came up with the causal/historical theories in the form of the “Madagascar” or “Santa Claus” examples.³¹⁶ Just as we want to say that what we now mean is dependent upon our past use, but leave room for the possibility that we occasionally break free from it, we want to say that what we mean can depend upon our future use, but leave room for the possibility of its breaking loose from us.

Objection: If a community were to split, with each sub-community developing accessible but incompatible equilibria from the original community’s use, then it would seem that we would have to attribute both of the incompatible meanings to the original community’s utterances.

Reply: Our entitlement to read our own development back into past use comes from our predecessors and ourselves being part of a single, developing, linguistic practice. Because of this, one can account for such cases by insisting that linguistic practices, like much else, do not preserve their identity through splits. The fact that, when such splits occur, the existence of a rival community undermines each community’s entitlement to say what their predecessors meant does not entail that such entitlement is not present when no splitting occurs. An analogous point is frequently made with respect to personal identity: the fact that, if I were to split, I could not be identified with either of my successors does not, in itself, entail that I cannot be identified with my ‘future self’ when I don’t split. An *actual* split undermines any identity between past and future, but the mere *possibility* of a split need not do so.

³¹⁶ See Evans 1973, Kripke 1972.

Objection: T-externalism still seems to entail at least that, when such splits occur, nothing determinate is originally meant by the term in question.

Reply: This objection is something we should take in stride. There are, after all, structurally similar cases within Burge's synchronic social picture and causal/historical accounts of proper names. That a community could split is no more a problem for the T-externalist than the possibility that a speaker's use of a name might be historically connected to two people is a problem for Kripke,³¹⁷ or the possibility that a speaker might unknowingly be part of two linguistically divergent sub-communities is a problem for Burge.³¹⁸ These possibilities are not *counterexamples* to the accounts in question; rather, they just suggest that, in some special cases, the accounts entail that nothing determinate will be meant by a term. This would be a problem if our intuitions suggested that something determinate was meant in these cases, but our intuitions about these cases tend to agree with the predictions of the accounts.

The intuition that we must *always* have thoughts with a determinate content must be given up if one is sympathetic with the causal/historical or social theories. The possibility of a split future thus does not present a novel type of problem that the T-externalist must face -- it is an instance of a phenomenon that all externalist (perhaps even all extensional) accounts must admit. Indeed, just as it should be viewed as a problem for a theory if it *did* come up with a determinate referent in the "Napoleon" case, it is a virtue of the above account that it predicts that there will be no determinate extensions in such split future cases.

³¹⁷ See, for instance, Evans' discussion (Evans 1973) of the possibility of our information about "Napoleon" actually coming from two different sources (one of whom did everything up until the Italian Campaign, and another of whom fought at Waterloo).

³¹⁸ Stich tells a story of this sort relating to the British and American use of "endive" and "chicory." The extensions of the two terms are reversed in the two countries, and so a speaker who spent an equal amount of time in each community, and felt the pull of both sets of norms, would fail to mean anything determinate by either term. (Stich 1983, 63.)

Objection: Another worry relates to the possibility of the use of the term having no future at all. If the Druids died out before the planes came by, the meaning of “ave” would remain indeterminate between *flying thing* and *bird*, since there would be no future linguistic behavior to make the distinction between them. The possibility of the term’s having no future use need not, of course, be so apocalyptic. The term might simply fall out of use and be replaced by another, or undergo a change in meaning. (If their term “ave” gradually came to stand for just pillows, then their reactions to the first planes would no longer be relevant.) As a result, T-externalism potentially leaves many of our utterances without any determinate content.

Reply: In such a case, one should (once again) simply admit that the meaning of the term in question is indeterminate. After all, the suggestion that what Edwin meant by “ave” remains indeterminate if the “ave”-using practice died out before planes were encountered should match our pre-theoretical intuitions about such cases. Furthermore, since all the alternatives to T-externalism leave what Edwin’s term refers to indeterminate in these cases as well, this can hardly be viewed as a reason for rejecting T-externalism in favor of one of these alternatives.

In any case, the increased visibility of this sort of indeterminacy cannot be held against T-externalism itself. Adding the community’s future linguistic development can only *add* to the determinateness of the contents in question, so any scenario where contents remain indeterminate for the T-externalist will also be indeterminate according to more temporally bound versions of externalism. The T-externalist framework may make this type of underdetermination more *visible* but it can hardly be blamed for the *existence* of such underdetermination.

Objection: T-externalism seems to tie our meaning anything determinate to our community’s having a non-branching future, but the determinacy of what we mean frequently seems independent of the future of our linguistic practice. For instance, by

1750, our ancestors would have meant *gold* by ‘gold’ even if they had all died out before they acquired the discriminating knowledge needed to pick it out as accurately as we do.

Reply: Of course it may be the case that the meaning of many, if not most, of our terms is independent of how things turn out in the future. However, this only shows that, for these terms, there is only one accessible equilibrium. As explained in section 4.1, a term’s extension can be “settled” at a moment if the term has the same extension in every history passing through that moment. The meanings of many of our terms (those for which there is a single accessible equilibrium) may be settled in this fashion. A term with an unsettled extension, however, should not be understood as having an ‘indeterminate’ extension, since having an indeterminate extension is itself something which can be settled. According to the T-externalist, a term with an unsettled extension could still have a determinate extension in *every* history, while a term with an indeterminate extension will not have a determinate extension in *any* history.

Our use of “gold” in 1750 may have been such a case of a term’s extension being settled, while the Druids’ use of “ave” manifestly is not. By 1750, the practice in which our use of terms like “gold” was embedded was explicit and developed enough to make a phenomenological or functional interpretation of “gold” inaccessible. The Druids could not, by hypothesis, have such a well-developed normative structure in place governing their use of “ave.”

It seems, then, that many of the reasons for thinking that there could not be a T-externalistic equilibrium for our semantic terms are not, in the last analysis, very compelling. The ultimate accessibility of any T-externalistic equilibrium will, then, depend upon the comparative accessibility of possible equilibria that take the contents of our utterances always to be settled at the moments of their utterance. Such accounts are the subject of the next section.

7 Non-temporalist interpretations of the phenomena

Any speaker's or community's usage up to a given time can underdetermine precisely what their terms refer to. Because of this, if one is to reject T-externalism and insist that all the relevant facts about an utterance's content must be in play by the moment of utterance, then one must claim that this underdetermination corresponds to an actual indeterminacy in what is meant. A non-T-externalistic semantics must thus be one that is tailored for indeterminate expressions. An obvious candidate of this sort is Field's partial reference based account of the semantics of languages with just such expressions.³¹⁹ Furthermore, while Field's account will be the focus of discussion in this chapter, but many of the criticisms that will be made of it (that it doesn't distinguish underdetermination from ontological confusion, that it fails to preserve continuity of meaning, etc.) can be carried over to other non-temporally based accounts such as those explaining such expressions in terms of degrees of truth. (None of the criticisms of Field's account have to do with his claim that certain utterances are neither true nor false.) How plausible we find T-externalism will depend to a large extent on how workable we find such alternatives.

While both the Druid and the Zebra cases are initially discussed by Wilson, Wilson himself suggests that the Druids' initial use of "ave" should be analyzed in terms of something like Field's notion of partial reference. According to the partial reference theorist, if, at a given time, the Druids' language (call it *D*) equally supports two rival interpretations, *I*₁ & *I*₂, that treat their term "ave" as referring to birds and flying things respectively, then "ave" *partially denotes* birds, and *partially denotes* flying things. Since the term partially denotes more than one (type of) thing, it doesn't *fully denote* anything. (A term fully denotes something if that thing is the only thing it partially denotes.)³²⁰ The interpretations *I*₁ & *I*₂ both *partially accord* with the semantics of *D*, since they both assign

³¹⁹ Field 1973, 1974.

³²⁰ Field 1974, 210.

to the words in *D* entities that the terms in *D* partially denote. Sentences in *D* can then be understood as true (false) if they are true (false) relative to *every* interpretation that partially accords with *D*. Sentences in *D* that have different truth values relative to different interpretations that partially accord with *D* are neither true nor false. On such an account, Edwin's sentence "Some aves are tasty" would be true, since it would be true on both *I*₁ & *I*₂. On the other hand, his sentence "Aves are found only on our island" would be false, since it is false on both interpretations. Finally, his sentence "All aves lay eggs" would lack a truth value, since it is true on *I*₁, and false on *I*₂.

Field uses his partial reference apparatus to explain Newton's use of the word "mass" (with "mass" partially denoting rest mass and partially denoting relativistic mass), and the T-externalist need not argue with this particular characterization. Indeed, this is a major difference between T-externalism and its more temporally bound competitors. The T-externalist can admit that there are some cases where, say, a partial reference account will be correct, but the partial reference theorist cannot allow cases where the T-externalistic semantic story is correct.

There is no denying that something like a partial reference account is appealing in some cases. For those cases where a partial reference account seems plausible, a T-externalistic semantics for cases of indeterminacy can be given that is structurally very similar to that proposed by Field.³²¹ In particular, one could give a partial reference semantics for those cases where a term's meaning is 'settled indeterminate.' A term's meaning will be settled indeterminate in precisely those cases where partial reference accounts will be most appealing. In particular, when the practice of using a word splits into two sub-practices (as was the case with "mass"), when the practice dies off, when no equilibrium is available for the entire practice, etc. The sentences that come out as true or false on Field's model would have the same values on such T-externalistic theories, and for those sentences that involve words whose extension is settled indeterminate, the T-externalist semantics could also be

³²¹ Or, if one prefers accounts that rely on degrees of truth to Field's account, one can build such mechanisms into one's T-externalistic semantics as well.

identical to the partial reference one. It is only *some* of the sentences that the partial reference theorists take to have no truth-value that the T-externalist need consider to be settled in the future.

On the other hand, while T-externalistic semantics can incorporate most of the mechanisms associated with Field's account, the converse is not the case. *All* of those terms whose extension the T-externalist can treat as "unsettled" the partial reference theorist must treat as having an indeterminate extension. (Compare the claim that statements about the future are neither true nor false to the claim that their truth value is unsettled.) Because of this, while various examples can be given where partial reference accounts can seem very plausible, the T-externalist can absorb such examples into his account. The debate must focus, then, on precisely those cases that the T-externalist appeals to, since it is only on these that there will be a difference between the respective judgments of the theorists in question. The question becomes, then, whether the class of terms whose extension is underdetermined by current use should be seen uniformly as having indeterminate extensions (as the partial reference theorist suggests), or whether they should be divided into two classes, settled indeterminate and unsettled (as the T-externalist suggests).³²²

Indeterminacy has almost always been presented as a semantic version of the underdetermination of theory by evidence, and underdetermination is typically presented as the possibility that two rival theories could both be compatible with all the available data, with the added claim that, in the case of semantics, there can be no 'fact of the matter' beyond that which can be determined. For most of what follows, the phenomena in question will simply be discussed in terms of underdetermination, since whether underdetermination at a time is actually indeterminacy is precisely with at issue between the T-externalist and someone who thinks that the facts available at the time are all the relevant facts.

³²² While Field originally presents the notion as one having to do with scientific terms and scientific revolutions, not terms in general, he seems to imply that his solution applies to all cases of indeterminacy.

How one feels about treating all cases of underdetermination as cases of indeterminacy will depend, to a certain extent, upon how important one takes the differences between the possible types of underdetermination to be. Consider, for instance, the following three possible types of underdetermination facing an interpreter.³²³

1. Underdetermination in which the extension of one of the rival concepts completely encompasses that of the other(s).
2. Underdetermination in which the rival concepts, while not coextensive, coincide (largely or) completely in the speaker's environment.
3. Underdetermination in which there is no overlap at all between the possible referents.

Underdetermination of the first sort can be seen in the initial use of "Grant's Zebra," and the second sort can be seen in Edwin's initial use of "ave." Cases of the third sort, however, are completely different. One is faced not with a common source of information that can be conceptualized in two different ways, but, rather, two distinct sources of information. This sort of underdetermination arises when, say, a speaker uses the term "Peter" for what he takes to be a single person, when 'Peter' is actually a pair of twin brothers, each of whom encounters the speaker equally often. Cases of the first two sorts should thus be distinguished from cases of the third, where the interpretee is 'ontologically confused.'³²⁴ An interpretee who is ontologically confused is very different from one whose use merely underdetermines what he means by his terms. While semantic underdetermination should be understood as occurring when two or more rival semantic theories can both be compatible with all the available semantic evidence, cases of ontological confusion clearly aren't instances of this. It's not that *both* theories (i.e.:

³²³ This list is not meant to be exhaustive. In particular, the underdetermination discussed here could be referred to as "practical underdetermination." There is no disagreement about what particular things the speakers are responding to, only about how those things should be classified. There is also, of course, "Quinean" underdetermination, which claims that the question of which particular objects the speakers are responding to will be underdetermined by the observational evidence (how can we tell they aren't responding to temporal slices of bird/flying-things, etc.). This latter purported type of underdetermination will not be discussed here (though for a further discussion of it, see Jackman, forthcoming).

³²⁴ For the use of this term to describe cases of type 3, see Camp 1987, 21.

“Peter” refers to twin #1 and “Peter” refers to twin #2)³²⁵ fit all the available semantic evidence. Rather it’s that *neither* does. Picking between the two twins would be like choosing between two rival physical theories each of which can only explain half of the phenomena. The evidence does not, admittedly, allow us to choose between the two, but this is because it determines that neither alternative is empirically adequate.

There are thus serious disanalogies between cases of ontological confusion and cases of underdetermination in physical theory. As a result, such cases are not best understood as being instances of semantic underdetermination. In terms of the apparatus from chapter one, what distinguishes semantic underdetermination from ontological confusion is that in the former case there are *multiple* equilibria available to make sense of the interpretee, while in the later there are *no* such equilibria available. Characterized in this fashion, Newton’s use of “mass” is a candidate for being understood as a case of ontological confusion. If Newton’s belief that momentum was the product of mass and velocity, and his belief that the mass of an object did not change with respect to its frame of reference, were both deeply enough entrenched, there would be no single equilibrium available for his use of “mass,” and some form of bifurcation would have been necessary. As a result, while rest and relativistic mass ‘overlap’ at low speeds in a way similar to how birds and flying things overlap on Edwin’s island, Newton may best be understood as ontologically confused in a way that Edwin is not because he cannot give up either of his conflicting general beliefs in the casual fashion Edwin can give up his belief that, say, all aves are living things.³²⁶ As a result, even if something like partial reference is the best way to deal with Field’s example, one should not assume that it is a the best way to deal with underdetermination in general. Indeed, it would be preferable to give underdetermination a semantic treatment that is

³²⁵ And these certainly aren’t the only two available. One could also imagine theories which cashed out the confusion in their account of the speaker’s predicates rather than in the satisfaction clauses associated with his singular terms.

³²⁶ Ultimately, not much hangs on the correctness of the characterization of Newton as ontologically confused. The fact that the subsequent practice split is enough to legitimate Field’s ascriptions in the T-externalist’s eyes.

distinct from ontological confusion. The T-externalist is able to do this, the partial reference theorist is not.

Furthermore, it should also be noted that partial-reference theories are not especially hospitable to self-knowledge either. Explaining the Wilson cases in terms of partial reference commits one to saying that we change the meaning of our term each time the extension is made more precise and one of its partial denotations is eliminated. This is a very unattractive feature of partial reference accounts, since (as stressed in chapter three) a theory of meaning should, if possible, be phenomenologically faithful to our assumption that the meanings of our terms stay fairly constant over time.³²⁷ Since the speakers often do not perceive themselves as changing what they mean by their words, a theory that says that they do must claim that they lack self-knowledge to at least that extent.³²⁸

Finally, it should be noted that all the considerations mentioned earlier in favor of the T-externalist account (increased determinacy, preservation of meaning over time, etc.) count against the partial reference one. These considerations are not, of course, decisive reasons for rejecting partial reference based accounts, but making room for a T-externalistic equilibrium need not involve showing that there are no alternatives available. Rather, it only requires showing that the T-externalistic equilibrium is palatable and that its competitors are not *obviously* preferable. Since the cases in dispute arise when a presupposition about language breaks down, it may be that any candidate equilibrium will

³²⁷ The appeal of this intuition is not lost on Wilson, who suggests that meaning be held constant during a term's denotational refinement. Wilson states that while the extension of a term can change when we learn more about the world, there is no reason to think that the meaning has thereby changed as well, and he justifies this by claiming that:

From the druids' point of view, the meaning of 'bird' has not altered, nor is there reason for the field linguist to describe matters in this fashion either. Attributions of change of meaning seem best tied to *recognition* by the community of a need for a conventional decision, and *ex hypothesi*, this is not the case here. (Wilson 1982, 572.)

³²⁸ Though, once again, there must be room for this to occur in cases of prolonged mistaken identity turning into change of meaning.

pinch in some places. As was discussed in the first chapter, we have no guarantee that there will be a single right answer about what equilibrium to reach in such disputes. Just as Linnaeus and Ishmael can disagree about whether or not to call whales are “fish,” without either party being guilty of any cognitive failing. The T-externalist and partial reference theorist could disagree over what to say about Edwin without either of them being mistaken. The purpose of this chapter has not been to show that one *must* adopt a T-externalistic equilibrium, but rather to show that if there are (multiple) equilibria available for our semantic terms, the T-externalistic equilibrium will be one of them.

8. Conclusion

It was suggested in the first chapter that disputes about meaning and content could be understood as the product of our trying to find an equilibrium between the general beliefs we associate with our semantic terms and the ascriptions we make when we attribute thoughts and sayings to others. The second and third chapter stressed the importance of our self-interpretations when reaching such an equilibrium, and argued that our practice of interpreting others in terms of social usage could be incorporated into an equilibrium because speakers understand *themselves* as engaging in a shared practice when using language. In much the same way, this chapter has argued that we can endorse everyday ascriptions of content that reflect a sensitivity to future usage because we understand this shared practice as itself extending through time. Most of the potentially unintuitive consequences of allowing future use to contribute to what we mean were dealt with on one of two ways. First of all, by distinguishing possible from accessible equilibria in the way suggested in chapter one, many of the consequences that might have been associated with radical shifts or splits in usage disappear. Secondly, many of the consequences (such as, say, those associated with self-knowledge and behavioral explanation) that come from adopting a T-externalistic position have already been accepted in virtue of incorporating the socially infected ascriptions Burge appeals to. Furthermore, the ways in which those

consequences were made palatable for the non-individualist discussed in chapter three work equally well for the T-externalist. Alternative ways of describing the Wilson cases were discussed and found to have a considerable number of unintuitive consequences of their own. Finally, it was argued that since attempts to find a more temporally-bound equilibrium for such cases turn out to be no more intuitive than their T-externalistic competitors, there is no reason why we cannot endorse our temporally sensitive ascriptions in these cases.