

THE NONLOCALITY OF MENTAL CONTENT

Brian Epstein
Department of Philosophy, Virginia Tech
223 Major Williams Hall
Blacksburg, VA 24060
bepstein@vt.edu

INTRODUCTION

In 1881, the British physicist James Ewing observed a peculiar property that iron exhibits under the repeated application of magnetic fields. A bar of iron can be magnetized over time by applying a magnetic field to it. And it can be demagnetized by applying an opposite field. But if, having done that, the positive field is once again applied to the same bar, the second time it magnetizes faster than the first. It is as if it remembers having once been magnetized, even after its magnetism has been erased. To predict the magnetization curve of a piece of iron, it is necessary to look not just at its magnetization *state*, but at its magnetization *history*. Ewing named this phenomenon *hysteresis*, after the Greek word meaning “to come later.”¹

In physics, hysteresis is not indispensable. The reason a piece of never-magnetized iron has different characteristics from a once-magnetized piece is that there is a physical difference between the two, even if it is hidden from straightforward measurement. Rather, the past is an explanatory factor only out of convenience: it is not that the distant past actually has a direct impact on the present.

Traditionally, language and thought were taken to be similar to physics, with respect to their relation to the past. While we do speak of or think about things in the past, the past itself was thought to be only a crutch or way-station in the explanation of content. Historical or diachronic factors may be of use in accounting for concepts of historical objects, but on the traditional view these do not figure into the individuation of thoughts. Since the advent of externalism with respect to mental content, however, it has been evident that this is not likely to be correct. Although the arguments of Putnam, Burge, and others predominantly focus on the role of environmental and social factors in individuating mental content, from the outset it was also clear that historical factors, not only features of the synchronic environment, must figure into the individuation of thoughts as well.² A number of philosophers have conveyed this point by noting that even something as simple as a footprint is not individuated by its synchronic characteristics. For something to be a footprint requires that the past have been a certain way: a mark is a footprint only if

¹ See also John Elster, "A Note on Hysteresis in the Social Sciences," *Synthese* 33 (1976).

² Tyler Burge, "Cartesian Error and the Objectivity of Perception," in *Subject, Thought, and Context*, ed. P. Pettit and J. McDowell (Oxford: Oxford University Press, 1986), Tyler Burge, "Individualism and Psychology," *Philosophical Review* 125 (1986), Tyler Burge, "Individualism and the Mental," *Midwest Studies in Philosophy* 4 (1979), Tyler Burge, "Intellectual Norms and Foundations of Mind," *Journal of Philosophy* 83 (1986), Tyler Burge, "Other Bodies," in *Thought and Object: Essays on Intentionality*, ed. Andrew Woodfield (Oxford: Clarendon Pr, 1982), Tyler Burge, "Wherein Is Language Social?," in *Reflections on Chomsky*, ed. Alexander George (Oxford: Blackwell, 1989), Hilary Putnam, "The Meaning of 'Meaning'," in *Philosophical Papers* (Cambridge: Cambridge University Press, 1975).

it was struck in the past by a foot.³ Analogously, concepts do not only come to be through empirical connections made in the past, but the historical features of the world with which connections were made continue to figure into their individuation.

In some ways, the diachronic individuation of thought is even more radical than the falsity of individualism, since it means that hysteresis in semantics, unlike in physics, cannot be dispensed with. One way to articulate this point is by contrasting different failures of what we might call *locality*. The rejection of the individualistic individuation of mental content is a denial of the spatial locality of a thought to its conceiver. The indispensability of hysteresis involves a failure of temporal locality as well. In looking for the factors involved in individuating a thought, temporally and spatially nonlocal factors must be taken into account as well as local ones.

Even as individualism with respect to thought contents has been rejected, however, a weaker sort of locality principle has come to take its place in prevailing semantic theories. It is a constraint that admits, but closely circumscribes, the role of the past in the individuation of concepts. This revised constraint is implicit in a widely-held picture of how the past gets involved in thoughts in the first place. Surprising as anti-individualism is, there is wide consensus as to its basic explanation. It arises from the fact that making empirical connections with the world plays a critical role in constituting semantic content. David Wiggins has an interesting discussion of this in *Sameness and Substance*.⁴ Wiggins notes that Putnam's treatment of concepts of natural kinds resembles a number of historical accounts that oppose the empiricist tradition of accounting for natural kinds in terms of nominal essence. Putnam agrees with these accounts that concepts of natural kinds are not individuated in terms of manifest properties and relations or appearances. Also in common with these accounts is the claim that the basis for individuating the kind is some physical or environmental principle. It is not the rejection of this tradition that is distinctive about Putnam's account. What is distinctive, observes Wiggins, is explicitly tying the determination of semantic content to the *demonstration* of a sample in the environment, when a concept is introduced. It is the connection to the world that grounds the empirical applicability of concepts. This point is compatible with the important insight that we, even as thinkers, are part of the world, not disembodied recipients of stimuli from the world. Absent empirical connection, mental states would fail to be intentional.⁵

Most empirical connections, of course, are not occurring presently, but took place in the past. Where the past enters semantics, this suggests, is in the preservation over time of those factors introduced in virtue of the empirical connections made with them. History becomes a factor as empirical connections recede into the past. The factors to which we make empirical connections, together with the passage of time, is thus seen to account for instances of hysteresis in mental content. A person who lived in the past, or a sample of a substance encountered in the past, may thereby figure into the individuation of a concept that is employed today. Such a person or sample might be epistemically inaccessible to us today, but nonetheless is part of the network of members of the linguistic community as a whole, and the

³ Cf. Fred I. Dretske, *Explaining Behavior : Reasons in a World of Causes* (Cambridge, Mass.: MIT Press, 1988), Robert Stalnaker, "On What's in the Head," *Philosophical Perspectives* 3 (1989).

⁴ David Wiggins, *Sameness and Substance* (Cambridge, Mass.: Harvard University Press, 1980), 87.

⁵ This is related to a point first made by Tyler Burge in connection with the primacy of empirical engagement in the possession of *de re* attitudes. Cf. Tyler Burge, "Belief De Re," *Journal of Philosophy* 74 (1977): 346.

parts of the world with which the community has had contact. Though these factors in the determination of semantic content are not local to a contemporary speaker, they remain local to the community. The picture that emerges is of a network of transmitted concepts, with empirical connections at the nodes, and the discursive community living on, thinking, acting, and communicating – all of which are eligible to figure into the determination of thought content.

With this picture of the determination of concepts, then, comes what we might call a “weak locality” principle: *factors that figure into concept individuation must have an empirical connection to the discursive community*. Inasmuch as any external, social, or historical factor figures into concept individuation, it is in virtue of its connection to the community. The weak locality constraint is implicit in such approaches as causal and informational accounts of mental content, and analogous constraints are implicit as well in causal and anaphoric theories of reference.⁶

My aim in this paper is to challenge this principle. I think it has not been seen, even by a number of people who endorse Putnam/Burge-style arguments, that a variation on these arguments themselves rules out weak locality. And hence it has not been seen that accounts of concept individuation which assume weak locality cannot be sufficient. If externalism is correct, then more radical implications follow than many philosophers have supposed.

I will argue against weak locality with a new set of variations on Putnam/Burge-style cases. Burge, in one of his discussions of anti-individualism, usefully points out the role of external examples in individuation. However, his treatment of the connection between the corrigibility through testing of our characterization of concepts and the role of examples in individuation incorrectly suggests that they are limited to those with which we can be empirically connected. Integrating semantic hysteresis with the use of examples in individuation demonstrates that we cannot exclude nonlocal factors from concept individuation. I will present two preliminary sets of cases: first I will clarify aspects of how the past figures in concept individuation, and then I will reconsider Burge’s demonstration as to how environmental examples do so as well. These provide the groundwork for the new approach, demonstrating that semantic hysteresis occurs even with historical factors having no causal connection to the discursive community.

HYSTERESIS IN DETERMINING CONTENT

I will begin by pointing out some aspects of the role of the past in individuation. A familiar example from Putnam, illustrating the dependence of semantic facts on the past, is his discussion of “robot-cats.”⁷ Putnam imagines that, unbeknownst to us, 50 years ago Martians came to earth and replaced all the cats on the planet with robots, which are currently indistinguishable to us from biological cats. He then asks what, in these circumstances, the truth-value of the sentence “Cats are animals” is. There once were cats in the environment, Putnam points out, and now the Martians have replaced cats with something else. But in replacing them, the Martians have not falsified the generalization that all cats

⁶ Many teleological views of semantic content also adhere either to this weak locality thesis, or to a related and somewhat weaker locality thesis. My arguments won’t directly address this weaker thesis, but I believe the same arguments to be extensible to a number of teleosemantic approaches.

⁷ Putnam introduced this example in Hilary Putnam, “It Ain’t Necessarily So,” in *Philosophical Papers* (Cambridge: Cambridge University Press, 1962), Putnam, “The Meaning of ‘Meaning’.”

are animals, even though we cannot distinguish the new stand-ins from cats.

This example, if cast properly, demonstrates hysteresis with respect to semantic content.⁸ As Putnam sets up the thought-experiment, the actual and the counterfactual situations are not synchronically indiscernible. To put the situation in the form of a standard nonsupervenience argument, it is useful to reverse the history. Consider the counterfactual situation in which cats were historically robots, and 50 years ago were replaced with biological organisms identical to our cats. In the actual situation, the thought that *all cats are animals* is true. In the counterfactual situation, however, it is false. With the replacement, the Martians eliminated robots from the earth, and their simulacra are biological organisms, not what cats counterfactually were.⁹ The thoughts that the individual is having in the actual situation are ones about biological organisms, and in the counterfactual situation are about robot-things. Nonetheless, the situations are synchronically indiscernible: in both the counterfactual and the actual situations, the things on the planet today are biological organisms. Thus the thoughts do not supervene on the synchronic state of the world. This case, and ones like it, demonstrate not only anti-individualism, but that historical factors figure into concept individuation. Here, hysteresis arises from the diachronic preservation of content. *Cat* might be set up at a single point in time, or it might be set up with a pattern of historical usage or reproduction, and its use in the present recalls those factors.

However, semantic hysteresis does not only arise as empirical connections fade into the past. Rather, the past can be involved in concept individuation as soon as a concept is first possessed. As a preliminary, consider the standard role of external factors in name introduction. Suppose, for instance, that I intend to assign a name to a species of the cat family, using a newborn animal. Not being a zoologist, I do not know what kind of cat it is: suppose I have been given one of either a cheetah or a lynx, but I do not know which. To the species of which this newborn is a member I assign the name ‘catkind 1’.¹⁰ It does not require that I be able to discern the difference, in order for the dubbing to be successful: if the sample happens to be a newborn cheetah, then ‘catkind 1’ refers to the cheetah; if a newborn lynx, then it refers to the lynx. The case depends on the fact that newborn cheetahs and lynxes

⁸ The general structure of demonstrating nonsupervenience is well-known. It involves describing two situations, an actual one and a counterfactual one, in both of which an individual has a particular thought. The two situations are identical in every way with regard to the internal state of the speaker’s mind, excluding intensional properties. The counterfactual one differs from the actual one in a particular respect, typically involving the speaker’s environment or social context. It is shown that the thought, which is true in one situation, is false in the other. This implies that the determinants of content are different in the two situations, and hence don’t involve the state of the speaker’s mind alone. My principal intent here is to build off of these arguments, rather than to answer potential objections to the standard arguments.

⁹ In the counterfactual situation, ‘cat’ may be a non-natural kind term, since robots are produced by the Martians. Depending on one’s interpretation of the semantics of non-natural kind terms, it could be that biological cats are encompassed by this category, so that *some cats are animals* is currently true in the counterfactual situation. Even so, *all cats are animals* will be false, since cats also encompass robots. It is also possible to recast the case with the Martians replacing one biological kind with another, to make the claims whose truth value changes more general. Another solution is to reverse the situation somewhat differently, with two different counterfactual situations, where in the first, cats are robots and have been all along, and in the other, cats were once biological organisms but were switched to robots.

¹⁰ In doing so, we assume whatever mechanism for fixing reference that suffices to pick out the species, such as the explicit use of sortals in indicating the aspect of the sample to be picked out, so long as it’s short of giving a description that singles out one species uniquely on its own.

are indistinguishable to me, as the dubber of the term, and so the individuation of the referent is a matter of things external to the descriptive content I possess. In this case, the factors determining its reference are all synchronic.

The same process, though, yields terms that depend on diachronic factors, immediately upon their introduction. This arises in cases where two kinds being dubbed are synchronically indistinguishable, but where non-synchronic factors distinguish them. Most contemporary philosophers of biology, for instance, take species to be individuated in part in terms of the history of their development. Two organisms, then, can be identical in structure and yet belong to different species, if they have different lineages. Consider then two species of cat, one in Britain that descended from one ancestor species, and one from France that descended from a different ancestor species. Within each species, there is variation among individuals: because of this variation, suppose that one of the British cats is atom-for-atom identical to one of the French cats. Assuming species are not individuated based on the structure of individual members of the species alone, the British cat is a member of the first species and not of the second, and vice versa for the French cat, even though they are molecularly identical.

To apply the nonsupervenience test to thoughts about these species, suppose that we have two cats in the lab. In the actual world, cats A and B are of the British lineage, and in the counterfactual case cat A is French and cat B is British. We use cat A as the sample on the basis of which the species concept *catkind 2* is generated, and then subsequently have the thought that cat B is a *catkind 2*. In the actual case, the thought is true, and in the counterfactual case, it is false. The synchronic facts are identical in both cases, and so appeal to synchronic factors alone is inadequate for determining the thought content.

Ordinary artifacts, social, and cultural objects, are also not individuated by each synchronic snapshot of them. Consider a piece of music. One of the notes near the beginning of Debussy's "Prelude to the Afternoon of a Faun," for instance, is an unaccompanied A-sharp; and it happens that the same note occurs toward the middle of Madonna's "Like a Virgin." That the two coincide for some part of them, however, does not mean that they are the same piece.¹¹ Suppose we compose a new piece, consisting of the first half of "Like a Virgin" followed by the last half of "Afternoon of a Faun" (call it "Like a Faun"). This piece also is not the same as either the Debussy or the Madonna. To construct the standard argument, suppose in the actual situation that someone walks into a New York Philharmonic concert of "Afternoon of a Faun" towards the end, and, hearing the last few minutes of a piece being played, calls it 'Piece 1'. She then has the thought that Piece 1 was composed in 1894. The counterfactual situation is synchronically indistinguishable, except that instead of "Afternoon of a Faun," the listener has walked in on the end of a concert of "Like a Faun."¹² In the actual case, the thought is true, and in the

¹¹ Individuating pieces is a complex matter: a piece may retain its identity even when it is transposed, so that no notes coincide between performed versions; and it may also be retained in an improvisation or interpretation, as with improvised performances of jazz standards. But however we individuate pieces, a short overlap doesn't suffice: "Like a Virgin" isn't the same piece as "Prelude to the Afternoon of a Faun." The role of the past in individuating pieces, here, is independent of many considerations of the role of history in the individuation of a piece of music, such as whether it depends on historical aspects such as who the composer was. Cf. Peter Kivy, "Platonism in Music: A Kind of Defense," *Grazer Philosophische Studien* 19 (1983), Jerrold Levinson, "What a Musical Work Is," *Journal of Philosophy* 77 (1980).

¹² As told, the case isn't likely to be entirely indistinguishable synchronically: they'll have a memory of what they've played, etc. The case can be easily modified, though, to make the synchronic facts entirely indistinguishable in the actual and counterfactual cases.

counterfactual case, it is false, even though the situations again are synchronically indistinguishable. The thought, then, is not individuated by the synchronic facts alone.

These cases, so far, all satisfy a weak locality constraint. That thoughts about a species fail to supervene on the present is simply a consequence of species being temporally extended entities. It is less a matter of the nature of the thought than the nature of the object that is thought about.

EXAMPLES AS FACTORS IN INDIVIDUATION

The second aspect of anti-individualism I want to draw on is the role of examples as factors in individuation. The use of examples is important, because it indicates a more significant sort of dependence on historical factors than those mentioned above. Dependence on historical factors will involve neither the preservation of content nor the persistence of individuating factors through time. Rather, diachronic examples are a straightforward way to demonstrate that mental content depends on an indeterminately large number of nonlocal factors, factors in the past that need not be empirically connected with the community.

This point is a subtle refinement of the implications of Putnam's Twin-Earth case, but is crucial for understanding the role of the environment in concept individuation. Burge discusses the role of examples in individuation in "Intellectual Norms and Foundations of Mind," and there are aspects of his argument I want to highlight and build on. At the same time, Burge's argument has some conspicuous weaknesses to shore up, and more importantly, I will show in the following section that Burge's excessive focus on the way examples can serve as correctives to normative characterizations of concepts suggests a mistaken limit to which examples are eligible to figure into concept individuation.

Nonsupervenience tests are successful for demonstrating anti-individualism if they can demonstrate that thought contents differ across situations, fixing the individualistic properties of the conceiver across the situations, but allowing anything else in the context to differ. Typically, the intention is to change exactly one external characteristic of the environment. So a successful nonsupervenience case can make two different points: it demonstrates that thought-contents are not individualistically individuated, and it singles out a kind of factor in the individuation of the thought contents. In Putnam's twin-earth case, he considers differences between situations set in 1750, so that the only difference between earth and twin-earth is the substance that rivers, lakes, etc., are composed of. If, then, there is any explanation at all as to what individuates thoughts, the substance is a factor in the individuation of water-thoughts.

We might, on the other hand, want to ask a more specific question about the factors involved in individuating Oscar's (or Twin-Oscar's) thoughts. Do the thoughts depend on the samples with which the term 'water' was first introduced? Or is it all the liquid in Oscar's environment, or some large subset of it? One thing we know it is not, by analogy from the robot-cat case, is the synchronic samples of water alone. But there is nothing about Putnam's case to indicate that anything more than the dubbed samples figure into the individuation of Oscar's thought. In particular, the Twin-Earth case does not demonstrate that a series of examples of water figure into the individuation of the concept *water*.

In "Intellectual Norms," Burge has two goals, namely to generalize his arguments against individualism to a very broad set of concepts, and also to demonstrate that social practices are not the

only nonindividualistic factor in individuating mental states.¹³ These two goals are somewhat at odds, since the demands of generality force him to choose a case that is susceptible to some difficult challenges, which weakens the argument for other nonindividualistic factors. It is not my intention to provide a defense for anti-individualist arguments in general, but it is worthwhile nonetheless to shore up the case. Certain of its shortcomings can be moderated by limiting its generality, and focusing strictly on the role of examples in individuation.

Burge's well-known argument begins with the claim that any normative characterizations of a concept can be challenged. For empirically applicable concepts, Burge argues, it is not necessary that any member of the linguistic community have the explicational resources to characterize them correctly. Even the consensus of the most competent speakers, in fact, has no assurance that it is true.¹⁴ The testing of a characterization, then, can be informative. To apply this, Burge constructs a version of the standard nonsupervenience argument, involving the concept *sofa*. In the actual situation, an individual has thoughts about sofas, and lives in a society in which sofas are objects employed for sitting (i.e., they are strong enough to hold up under being sat on, etc.), and people theorize that they are so-employed. The individual then develops a radically non-standard theory, that sofas are religious artifacts. He has reasons for believing this theory, but in fact it is mistaken, and he can later learn this by testing the examples of sofas and their uses. In the counterfactual situation, the individual is indistinguishable from the actual case, as is his theory that sofas are religious artifacts. In that society, however, the things he calls sofas actually are religious artifacts (i.e., they are not strong enough to hold up under being sat on, etc.), the social practices in which they are involved are religious ones, and people in general believe that they are so-employed. In the actual case, then, a thought about sofas – namely, that the non-standard theory holds of them – is false. And in the counterfactual case, it is true.

Based on the argument that one can doubt even necessary truths regarding a concept, Burge points out that one can possess a theory that is as non-standard as one wants. The situation that makes that non-standard theory true, then, ensures that what is being talked about is not the same subject with different characteristics, but rather is a different subject altogether. In this case, Burge argues that the actual individual is thinking about sofas, but the counterfactual one can only be taken to be thinking about something else, not about sofas at all. In the counterfactual case, the non-standard theory is true, and so the truth values of the thoughts involving the normative characterizations are different in the two cases. The individual, however, is indiscernible across the situations; hence, the thought content does not supervene on the state of the individual.

Two problems with this case are exacerbated by the choice of *sofa*. The first is a criticism that many have levied against this argument of Burge's, one that makes even some people who are persuaded by anti-individualism arguments take this case to be questionable. The issue has to do with the attribution to the individual in the actual case a thought about sofas, even when we are willing to report of that individual, "She theorizes that sofas are religious objects."¹⁵ In the actual case, we ascribe to the actual

¹³ Burge, "Intellectual Norms and Foundations of Mind," 707.

¹⁴ *Ibid.*: 706.

¹⁵ Cf. Kent Bach, "Burge's New Thought Experiment: Back to the Drawing Room," *Journal of Philosophy* 85 (1988), Tim Crane, "All the Difference in the World," *The Philosophical Quarterly* 41 (1991), Reinaldo Elugardo, "Burge on Content," *Philosophy and Phenomenological Research* 53, no. 2 (1993).

conceiver attitudes toward sofas, namely, the doubt that sofas are things typically sat on, and rather the belief that they are religious objects. The problem is whether in that attribution, it is appropriate to move from reporting this individual's belief using the term 'sofa', to inferring that sofas are the object of the belief. The individual's genuine doubt that sofas are things sat on, on this criticism, calls into question what the thing is, about which she genuinely has a doubt.

A second problem, which is more central to our purposes, is that the case Burge presents, even if the nonsupervenience point is accepted, does not isolate *examples* as the external factors in individuation, as opposed to other factors. In spite of the fact that Burge says that the case is designed to show that social practices are not the only nonindividualistic factor in concept individuation, the sofa case does not successfully bracket social practices. The counterfactual situation Burge describes is different in a number of ways from the actual one: not only are sofas potentially rickety, but the social practices involving them are different, and the social theory is also different. It may be that those differences suffice to yield the difference in truth-value of the non-standard theory. It is not clear, in fact, how the case could be set up such that the theory about religious artifacts could be true while isolating only differences in the examples, rather than practices as well.¹⁶

Both of these problems can be addressed, to a large extent, by choosing a case that is less generally applicable than the sofa case. For the case of sofas, it takes a radically nonstandard theory, and hence a radically different counterfactual situation, to make it plausible that the individuals in the actual and counterfactual situations have different thoughts. This is why it is critical to Burge to argue, at the outset, that normative characterizations can always be dubitable. Specifically to demonstrate the role of examples, however, does not require quite so strong a point on this front, but does require choosing a case in which the contribution of examples to individuation can be isolated from social practices.

Consider a nonsupervenience test involving a conceiver having thoughts involving the concept *fire*. To isolate examples from social practices, place this case in an era when people do not have any social practices or theories involving fire. Suppose the community simply sees fire from afar, before it is harnessed, in both the actual and counterfactual situations. In both situations, suppose the indiscernible conceivers have a theory of what fire is. Suppose they, plausibly enough, take it to be the shimmering afterglow of a spell. So A has a belief that we can characterize as:

(1) Fire is a magical afterglow.

Suppose this theory has testable implications: for instance, as with other magical phenomena, that the spell is broken and the fire disappears when one gets close to it. Conceiver A, in the actual situation, is mistaken. Testing her theory, she gets burned time after time, and so has to give it up. Conceiver B, in the counterfactual situation, is correct about her theory, and her test is correspondingly successful. The

¹⁶ In a footnote, Burge gives a more direct consideration for the necessity of examples. A's hypothesis is in part based on his claim that common perceptions regarding sofas are non-veridical. This is the source of his doubt about what he takes (correctly) to be the views that most people have, i.e., he hypothesizes that most people are mistaken in their understanding of the things they're coming in contact with, and their own behavior with respect to these things. Given the possibility of this, the communal theory alone can't be the arbiter of what sofas are. But still, it is arguable that it is the social practices, the misperception of which the society has led astray, that does the work in A's hypothesis. Someone doubting the indispensability of examples could agree that one can doubt a theory that's necessarily true of sofas, but deny that it is consistent to hold that the social practices carve out one thing and that the examples another, in this case. Burge, "Intellectual Norms and Foundations of Mind," 716-17.

glow disappears each time she walks towards it.

For the case of fire, it is plausible that one can have nonstandard theories involving it, without threatening the attribution problems that have dogged the sofa case. The difference between the actual and counterfactual cases is significant enough that, as in the basic Twin-Earth case, different normative characterizations will be true in the different circumstances. Nonetheless, it is as reasonable to attribute to conceiver A thoughts about fire as it is to attribute to Putnam's Oscar thoughts about water. That Aristotle, for instance, believed that water and fire were elements does not imply that that his concepts of water and fire were implicitly defined by his theory, rather than being individuated in large part by external factors. The corrigibility of conceiver A's theory, in the face of examples, indicates that the actual properties of the example figure into the individuation of the concept *fire*, as employed in her thought.

What is demonstrated in this case is subtly but critically different from the conclusion of Putnam's Twin-Earth case. As I mentioned, the Twin-Earth case does demonstrate that samples of water figure into the individuation of the concept *water*, but they leave it undecided as to the role that samples play. The fire case is less general than the sofa case, but it demonstrates that examples other than used in the initial introduction of the concept play a role in the concept's individuation.

Burge regards the corrective potential of examples as illuminating an aspect of the nature of thought. "Our conception of mind," he states, "is responsive to intellectual norms which provide the permanent possibility of challenge to any actual practices of individuals or communities that we could envisage."¹⁷ The properties of examples of fire enable us to challenge our normative characterization of fire. It is unclear from Burge's remarks, however, whether examples play a role in concept individuation inasmuch as they can serve as such intellectual norms; or rather, whether their serving as intellectual norms is evidence that they play a role in concept individuation, but does not explain or constitute the role they play in concept individuation. The former is at least a plausible interpretation of Burge's diagnosis, i.e., that the dialectic of testing is part of accounting for what concepts are, such that examples figure into their individuation.

I will not challenge this former diagnosis directly, but only note that it does not follow from the ability of examples to serve as correctives to normative characterizations. Where Burge sees intellectual norms as constitutive of our conception of mind, the role of examples as correctives merely demonstrates that examples are factors in conceptual individuation, but does not entail that their corrective potential is an explanation of the fact that they are such factors. It is equally possible that their corrective potential is simply an outcome of, rather than an explanation for, the fact that examples serve as factors in conceptual individuation. In the following section, I will combine the corrective role of examples with their occurrence in the past, to show that even examples that have no potential role in testing nonetheless serve in concept individuation.

PUNCTUATED FACTORS AND DISPENSABLE EVIDENCE

Considering examples in the past, we can see that factors in concept individuation go beyond those with which even the diachronic discursive community is connected. To demonstrate

¹⁷ Ibid.: 720.

nonsupervenience, consider a case in which the actual and counterfactual situations are synchronically indiscernible but different in terms of diachronic examples. The fire case can be modified to accomplish this, at the cost of introducing some artificiality. Consider a future situation in which fire has been eliminated, replaced perhaps by electric lights and induction cooking. Synchronically, then, there are no examples of fire. As with the previous case, suppose the conceiver has a belief about the nature of fire: that, for instance, it is magical. The differences between the situations are limited to examples in the past. In one situation, the history of the world is identical to the actual world, i.e., in which the examples are fire are the actual ones. Conceiver A, in this situation, has the belief characterized by (1). The only difference in the other situation is historical: the examples are identical to that of the earlier counterfactual situation, i.e., in which the examples satisfy the theory of magical afterglows.

I think it is straightforward to see that, as in the previous fire case, in the first situation the individual's thought is of fire, and in the counterfactual situation his thought is of a different stuff. But it is also worthwhile to give an argument for it, to strengthen confidence in this conclusion, as well as to begin clarifying the role that past examples have in individuation.

The reason it even needs an argument is that there is an important resource that the sofa case and the earlier fire case have, and which this one fails to: namely, a mechanism for testing the theory to evaluate it. In the above cases, it was demonstrated that examples are involved in individuation, by highlighting their role in the dialectic of refining characterizations of meaning. But with both situations synchronically indistinguishable, there is no basis for using examples in testing. The case stipulates that the extant examples are insufficient to decide between the theories, and so testing cannot decide between the cases. (Notice that this is the case for all cases of semantic hysteresis, including those involving just the preservation of content: no amount of testing will decide between the truth or falsity of the cat thoughts either, since they too are identical synchronically.)

The method for arguing this will be to re-introduce a synchronic difference between the situations, sufficient to let testing occur, and then to show that this difference is not the factor in the individuation of the thoughts, nor does its connection to the present make the past cases figure into individuation. To the case of future fire, add a synchronic difference: let there be documentary evidence in each of the situations – text, photographs, illustrations, designs – of the way fire once was. The respective theories, then, can be tested against the evidence. In the actual situation, the documentation denies the individual's theory, and in the counterfactual situation, it confirms it.

The documents are the synchronic things used in the dialectic of testing. But the reason that they are taken to be factors in the refinement of a normative characterization of *fire* is not that they themselves play a role in its individuation. Rather, it is that they provide evidence for the way fire was, and it is the former examples of fire that are factors in the term's individuation. It is not the evidence that has the normative force, but the fact that the evidence says something about the nature of the examples. The same occurs in any testing situation. We may have, for instance, only indirect tests that indicate that water is H₂O. It is not, however, either electrolytic reactions or the electron microscope or the light rays that individuate the concept, but the structure of water itself. Conversely, the role of examples in testing is a phenomenon of their role in individuation, but does not exhaust that role. In particular, where examples are not testable, that does not entail that they do not figure into individuation.

It is true that the documentary evidence is caused by former examples of fire with which the community was connected. But the documentary evidence is not only evidence as to what the structure of

those examples is, any more than some kind of evidence that smoking causes lung cancer among some sample population is just evidence for the causes of lung cancer *in that population*. The evidence indirectly implicates smoking for that subpopulation, and it equally indirectly implicates smoking for the population as a whole. Likewise, the documentary evidence of former fire, to be sure, is evidence for the structure of the examples that the documentarians encountered. But it is equally evidence for the structure of the enormous number of examples that were outside their experience.

The detail of these situations can obscure an intuitive point: inasmuch as examples are used in individuation, it is the exception to have the bulk of cases, even testable ones, in the present. And it is not testing or testability that makes even a testable factor play a role in individuation. Historical examples, though our contact with them is mediated through evidence, themselves are the things that figure into the dialectic of refining the characterization of concepts. This is apparent if explicitly historical things are considered, which can but do not often occur – such as, for instance, an infrequent political event like an abdication or a constitutional convention; some musical or dance form, like a tarantella or sarabande; an economic thing, like a royal bank or a ducat, etc. And the point applies equally to terms that are synchronically common as well, like fire or sofas. Some of the current examples of these can be involved in a dialectic of refining characterizations by testing them directly, and some of the past examples can be involved in this dialectic through indirect means; but this difference does not imply that the untested or untestable examples themselves fail to figure into its individuation. Though nonlocality is most easily demonstrated using hysteresis, it is also clear that it is not just diachronic factors that are nonlocal. Once the role of nonlocal diachronic factors is pointed out, using the dispensability of evidence, the same argument can be seen to apply to synchronic nonlocal factors as well.

The failure of weak locality implies that semantic content, to some extent, is hidden, or hermetic. On the traditional view of thought, semantic content is accessible to each conceiver, through introspection. On the prevailing externalist view, it is not accessible in fact, but still is accessible in principle. We may never learn facts about the local environment, in which a term is introduced, or may forget them as language is transmitted. But if we had always been perfect cognizers and scientific investigators, with perfect memories, on the prevailing view there is nothing that in principle would prevent us from knowing all there is to know about the contents of our thoughts. The nonlocality of content-determination, however, foils this. Even in principle, there are limits to the extent to which we can gain full mastery of our contents.

Any theory of semantic content relying on causal connections or the transmission of information, or one that takes thoughts to stand in anaphoric relations to empirical connections in the past, then, cannot be sufficient. Nonlocal factors figure into the individuation of thoughts; and so thought content cannot depend on networks of empirical connection alone. This does not negate the insight that empirical connection is critical for accounting for semantic content altogether. Empirical connections with the environment are surely necessary, for our environment to be ours. Those empirical connections, however, do not exhaust what counts as our environment. To determine what environmental factors are eligible to figure into the individuation of our thoughts, an account is needed of the individuation of the environment of a discursive community. The empirical connections of that community presumably figure into the individuation of what counts as the community's environment. But it is the community's environment as a whole, rather than the elements with which the community happens to have been connected, that can figure into the determination of thoughts in that community.

REFERENCES

- Bach, Kent. "Burge's New Thought Experiment: Back to the Drawing Room." *Journal of Philosophy* 85 (1988): 88-97.
- Burge, Tyler. "Belief De Re." *Journal of Philosophy* 74 (1977): 338-62.
- . "Cartesian Error and the Objectivity of Perception." In *Subject, Thought, and Context*, edited by P. Pettit and J. McDowell, 117-36. Oxford: Oxford University Press, 1986.
- . "Individualism and Psychology." *Philosophical Review* 125 (1986): 3-45.
- . "Individualism and the Mental." *Midwest Studies in Philosophy* 4 (1979): 73-121.
- . "Intellectual Norms and Foundations of Mind." *Journal of Philosophy* 83 (1986): 697-720.
- . "Other Bodies." In *Thought and Object: Essays on Intentionality*, edited by Andrew Woodfield, 97-120. Oxford: Clarendon Pr, 1982.
- . "Wherein Is Language Social?" In *Reflections on Chomsky*, edited by Alexander George, 175-91. Oxford: Blackwell, 1989.
- Crane, Tim. "All the Difference in the World." *The Philosophical Quarterly* 41 (1991): 1-25.
- Dretske, Fred I. *Explaining Behavior : Reasons in a World of Causes*. Cambridge, Mass.: MIT Press, 1988.
- Elster, John. "A Note on Hysteresis in the Social Sciences." *Synthese* 33 (1976): 371-91.
- Elugardo, Reinaldo. "Burge on Content." *Philosophy and Phenomenological Research* 53, no. 2 (1993): 367-84.
- Kivy, Peter. "Platonism in Music: A Kind of Defense." *Grazer Philosophische Studien* 19 (1983): 109-29.
- Levinson, Jerrold. "What a Musical Work Is." *Journal of Philosophy* 77 (1980): 5-28.
- Putnam, Hilary. "It Ain't Necessarily So." In *Philosophical Papers*. Cambridge: Cambridge University Press, 1962.
- . "The Meaning of 'Meaning'." In *Philosophical Papers*. Cambridge: Cambridge University Press, 1975.
- Stalnaker, Robert. "On What's in the Head." *Philosophical Perspectives* 3 (1989): 287-316.
- Wiggins, David. *Sameness and Substance*. Cambridge, Mass.: Harvard University Press, 1980.