Chapter 1

Introduction and Fundamentals

Abraham stood yet before the LORD. And Abraham drew near, and said, Wilt thou also destroy the righteous with the wicked? Peradventure there be fifty righteous within the city: wilt thou also destroy and not spare the place for the fifty righteous that are therein?.... And the LORD said, If I find in Sodom fifty righteous within the city, then I will spare all the place for their sakes. And Abraham answered and said.... Peradventure there shall lack five of the fifty righteous: wilt thou destroy all the city for lack of five? And he said, If I find there forty and five, I will not destroy it. And he spake unto him yet again, and said, Peradventure there shall be forty found there. And he said, I will not do it for forty's sake. And he said unto him, Oh let not the LORD be angry, and I will speak: Peradventure there shall thirty be found there. And he said, I will not do it, if I find thirty there. And he said, Behold now, I have taken upon me to speak unto the LORD: Peradventure there shall be twenty found there. And he said, I will not destroy it for twenty's sake. And he said, Oh let not the LORD be angry, and I will speak yet but this once: Peradventure ten shall be found there. And he said, I will not destroy it for ten's sake. And the LORD went his way, as soon as he had left communing with Abraham.

Genesis 18: 22–33
King James Bible

Perhaps the only point on which all theorists of vagueness agree is that vagueness is a form of unclarity—specifically, an unclarity about the boundaries of things. In language, vagueness concerns the extent of a term's application: There is no clear or definite boundary between the items to which the term applies and the items to which it does not. Some philosophers (e.g., Tye 1996, Hawley 2002, Maddy 2007, Barnes 2010) think that objects in the mind-independent, language-independent world also can be vague: Where does the mountain end and the valley begin? But the primary locus of vagueness is natural language, and linguistic vagueness is what I will be talking about.

'Tall', 'blue', 'heap', 'rich', and 'old' are prime examples of vague words: No clear line divides the tall people from the above average, or the blue objects from the green, or the old people from the middle-aged. In contrast, people who are precisely 6 feet tall are clearly divided from people of any other height, and objects that reflect light of precisely 475 nanometers are clearly divided from objects of any other reflectance. Most theorists of vagueness take the unclarity to be semantic, meaning that vague words lack sharp boundaries of application (they have "blurred" or "fuzzy" boundaries), but some take the unclarity to be epistemic, meaning that vague words have sharp boundaries whose locations are unknowable. I will say more about these different interpretations shortly.

As we start out, it is helpful to note how vagueness differs from several other species of linguistic unclarity, including underspecificity, ambiguity, and certain forms of indeterminacy. If you ask me when the meeting starts and I reply, "Between nine and five," my reply is precise, hence not vague, but it is *underspecific* for the purpose at hand. *Ambiguity* is the possession of two or more meanings, no matter whether vague or precise. For example, the ambiguous word 'odd' can mean 'peculiar' (presumably vague) or 'not evenly divisible

by 2' (presumably precise). The difference between vagueness and at least one species of *indeterminacy* is explained by Patrick Greenough:

> Suppose we stipulate that the open sentence '*x* is an oldster' is determinately true of every person sixty-eight years of age and over, determinately false of those persons sixty-five years of age and under, and neither determinately true nor determinately false of the remainder. If a speaker applies this term to persons who are between sixty-five and sixty-eight then we are entitled to say that they have done something not quite right and done something not quite wrong according to the dictates of the stipulation. But…'*x* is an oldster'…is *not* vague but rather, in some sense, semantically incomplete. This species of indeterminacy per se is not vagueness, since the term 'oldster' draws a perfectly sharp and clearly identifiable three-fold division across its associated dimension of comparison. (2003, 245)

Vagueness differs also from the kind of indeterminacy found in, say, the claim that Sherlock Holmes liked arugula. (Conan Doyle didn't say.) In principle this kind of unclarity could be removed without changing the meanings of any of the words involved. (Conan Doyle could have said.)[1] In contrast, vagueness is supposed to be a permanent unclarity whose removal would risk changing the meaning—indeed would threaten the very communicative utility—of a vague word. Crispin Wright asks us to

> suppose it is possible [to sharpen 'heap']. Then what in the semantics of [the word] is already inconsistent with our so refining [it]?…. 'Heap' is *essentially* a coarse predicate, whose application is a matter of rough and ready judgement…. It would for example be absurd to force the question of the execution of the

> command, 'Pour out a heap of sand here', to turn on a count of the grains. (1976, 333; emphasis added)

Some philosophers do conceive of vagueness as a form of indeterminacy, however. We will hear more about this idea.

Possession of unclear boundaries (semantic or epistemic) is often thought to be linked in some necessary way to two other linguistic phenomena: borderline cases and sorites paradoxes. Many theorists think that vagueness just consists in having borderline cases, where the latter are conceived as items to which it is unclear whether the word in question applies. More formally, borderline cases for a vague word 'Φ' are supposed to be *neither definitely (clearly) Φ nor definitely not-Φ*.[2] A cloth patch whose hue lies midway between a definite blue and a definite green may be neither definitely blue nor definitely not-blue, and a 65-year-old person may be neither definitely old nor definitely not-old. Christopher Kennedy and Louise McNally write that "the possibility of…'borderline cases' is one of the defining properties of vague predicates" (2005, 360), and Roy Sorensen notes "wide agreement that a term is vague to the extent that it has borderline cases" (2012).

Soriticality, the property of generating the notorious sorites paradox, is also often said to be criterial or constitutive of vagueness. For example, Otavio Bueno and Mark Colyvan claim that "a predicate is vague just in case it can be employed to generate a sorites argument" (2008, 5), and Wright asserts that "it would be inconsistent with elements already present in the semantics of [vague] predicates so to refine their senses that the sorites reasoning was blocked" (1976, 333). To see how the paradox gets going, suppose we are interested in the richness, measured by annual salary, of Americans aged forty to sixty in 2001. Then consider a series of salaries progressing from one that would make such an American clearly rich, say $200,000, to one that would make him clearly middle income, hence clearly not-rich,

say $50,000. (If you dislike these figures, feel free to substitute ones you prefer. The endpoints of the series need not be paradigm or prototypical cases—just clear or non-borderline ones.) Suppose further that each salary in the series is one dollar lower than its predecessor. Then it seems we can generate the following argument:

> $200,000 is a rich salary.
> For any number *n*, if $*n* is a rich salary, then $(*n* – 1) is a rich salary.
> Therefore $50,000 is a rich salary.

The premises seem true and the reasoning valid, yet we arrive at a contradiction: $50,000 both is and is not a rich salary.[3] Use of mathematical induction is not the culprit, for the absurd conclusion can be reached also by (among other things) a series of applications of *universal instantiation* and *modus ponens*:

> $200,000 is a rich salary.
> For any number *n*, if $*n* is a rich salary, then $(*n* – 1) is a rich salary.
> If $200,000 is a rich salary, then $199,999 is a rich salary.
> If $199,999 is a rich salary, then $199,998 is a rich salary.
> Etc.
> Therefore $50,000 is a rich salary.[4]

Soritical reasoning may be familiar from debates about abortion rights: Since a newborn infant is a person, and a human organism, say, one second younger than a person is also a person, it seems to follow that a conceptus is a person. The existence of the latter argument suggests that the word 'person' too is vague. (Of course, the conclusion that a conceptus is a person is not obviously absurd. Hence it is a matter of some urgency for defenders of abortion rights to figure out

what is wrong with soritical reasoning generally.) Dorothy Edgington observes that

> the paradox is not just a philosopher's puzzle, but something which affects our lives. There's the 'mañana paradox': the unwelcome task which needs to be done, but it's always a matter of indifference whether it's done today or tomorrow; the dieter's paradox: I don't care at all about the difference to my weight one chocolate will make. (1997, 296)

Theoretical difficulties caused by the paradox and by vagueness generally are not confined to philosophy and linguistics. Among other things, vagueness poses significant problems in the law (e.g., Waldron 1994, Endicott 2000, 2011), and in one way or another soritical reasoning threatens the coherence of the preference and indifference relations in rational choice theory (e.g., Quinn 1990, Ackerman 1994, Voorhoeve and Binmore 2006, Tenenbaum and Raffman 2012).

The major premise of the paradox (or each conditional premise) expresses the intuition that vague words *tolerate* incremental changes on the dimensions decisive of their application. As Wright explains, "there is with respect to any vague predicate the notion of a positive degree of change . . . insufficient to alter the justice with which [the predicate] is applied" (1975, 229). The mystery of the paradox is that, since any two adjacent items in a sorites series are only incrementally different, either both satisfy the predicate in question or neither does; either both salaries are rich or neither is. But then how is a transition made from 'rich' to any other predicate ('middle income,' 'borderline rich,' 'not-rich,' etc.)? How can there be—as surely there must be—a last rich salary in this apparently seamless, boundaryless, series?[5]

On its face, the paradox seems to show that many of our ordinary words are incoherent: Their use leads to contradictions. On the other hand, ordinary speakers use vague words all the time without landing

in absurdities like the one above; no competent speaker will keep applying 'rich' all the way down to $50,000. Indeed, far from paralyzing or incoherent, vagueness seems essential to our ability to communicate using natural language; vague words allow us to communicate easily, in a quick and casual way, without having to count grains of sand or dollars and cents of salary. These considerations suggest that the paradox is solvable, i.e., that it should be possible to discover what goes wrong in the paradoxical argument. Although several philosophers have argued that the sorites is incurable and vague words really are incoherent (e.g., Frege 1903, Dummett 1975), most have sought to exonerate vague language by clearing the puzzle away.

At a minimum, an adequate theory of vagueness needs to say exactly what vagueness is and what borderline cases are, and respond to the sorites paradox, and also make clear how these three phenomena—vagueness, borderline cases, and soriticality—are related. It has proven extremely difficult to do all of this; many attempts have been made. Since the project of this book is to provide such a theory, we need at the outset to have an idea of the major competing accounts that have been proposed and the kinds of difficulties they face. These other accounts have been amply reviewed and criticized elsewhere (e.g., Williamson 1994, Keefe and Smith 1999, Keefe 2000), so I will be brief.

1.1. WHIRLWIND TOUR OF COMPETING THEORIES OF VAGUENESS

Broadly speaking, we can distinguish four families of theories. Three of them take vagueness to be a semantic phenomenon, while the fourth construes it as epistemic.

(1) *Multivalued* approaches, including three-valued and degree theories, conceive of the blurred boundaries of vague words in terms of the

assignment of one or more values intermediate between true and false. According to three-valued theories, the sentence '$n is rich' is true of salaries that are definitely rich, false of salaries that are definitely not-rich, and *indefinite* of salaries that are borderline rich (viz., neither definitely rich nor definitely not-rich). Degree theories, which recognize infinitely many truth-values, hold that a borderline rich salary satisfies 'rich' to some degree intermediate between 0 and 1—say, 0.4 or 0.5.

Multivalued approaches offer a variety of solutions to the sorites paradox. For example, one three-valued account defines the conditional in such a way that the major premise is indefinite, hence untrue, because its instances are indefinite whenever the adjacent salaries at issue are borderline cases (e.g., Tye 1994). On this view the paradoxical argument is valid but unsound. According to one degree theory (e.g., Machina 1976), since each successive salary in the series satisfies 'rich' to a slightly lesser degree, each conditional premise is slightly less than true. However, the paradoxical argument is invalid because its conclusion, which is wholly false, is less true than its least true premise. Here *modus ponens* is less than fully valid. Another degree theory (e.g., Edgington 2001) has it that the paradoxical argument is valid because the 'unverity' of the conclusion is no greater than the sum of the unverities of the premises; essentially, the conclusion is no more false than all of the premises put together.[6]

(2) According to *supervaluationists* (e.g., Fine 1975, Keefe 2000), a sentence containing a vague term 'Φ' is true (false) just in case it is true (false) on every admissible way of making 'Φ' precise. More exactly: '*x* is Φ' is true *simpliciter* or definitely true just in case it is true on every complete admissible precisification of 'Φ.'[7] Roughly, a (complete) precisification is a way of making a term's application precise or "sharp" by classifying any unclear items either as Φ or as not-Φ; and an admissible precisification is one that also satisfies certain intuitive constraints. For instance, any admissible precisification of 'tall' will classify basketball player Shaquille O'Neal as tall (for a

human being), and if it classifies a given person as tall, it will also classify any taller person as tall. On this view, the major premise of the sorites paradox is false (*simpliciter*, definitely): Every complete admissible precisification of vague predicate 'Φ' establishes a sharp boundary between the extensions of 'Φ' and 'not-Φ'; hence the major premise is false on every such precisification. Borderline cases belong to the extension of 'Φ' on some but not all of its admissible precisifications; hence *x* is a borderline case for 'Φ' just in case '*x* is Φ' is neither definitely true nor definitely false. Although its semantics for vague words is gappy, supervaluationism is meant to preserve excluded middle: On each admissible precisification, every item either is Φ or is not-Φ—rich or not-rich, tall or not-tall—hence the sentence '*x* is Φ or *x* is not Φ' is definitely true, even where *x* is a borderline case. Thus the theory is not truth-functional: A disjunction can be definitely true even though neither of its disjuncts is. Of the four families of theories, supervaluationism enjoys perhaps the greatest following.

(3) *Contextualist* theories of vagueness take many forms, but their fundamental idea seems to be that vagueness is a form of context-sensitivity. According to one contextualist approach, a vague term like 'rich' is sensitive to contexts defined in part by the shifting verbal dispositions of competent speakers (e.g., Raffman 1994, 1996, 2005b). The most fully developed contextualist theory (Shapiro 2007) holds that vague expressions are sensitive to conversational contexts. On this dynamical approach, competent speakers have discretion to apply or withhold a vague predicate in borderline cases, depending on their conversational goals; borderline cases are items with respect to which competent speakers can 'go either way,' as Shapiro puts it (2007, 10). When a speaker chooses to apply a vague term to a borderline item, and his interlocutors acquiesce in his usage, a new conversational context or 'score' is established and the extension of the term is adjusted to include the item in question. In both "individual" and "conversational" versions of contextualism,

the sorites paradox dissolves essentially because there is no single context relative to which every instance of the major premise is true.[8]

(4) The fourth family of theories, *epistemic* theories, construe vagueness as a form of ignorance (e.g., Sorensen 1988, Williamson 1994, Graff 2000). Here the major premise of the sorites is false because there is in fact a sharp (though unstable) boundary between the rich salaries and the not-rich; it's just that we cannot know where the boundary lies.[9] Being a borderline case consists in being neither knowably rich nor knowably not-rich. According to one prominent version of epistemicism (Williamson 1994), we can't know where the boundary lies because its location is a function—also unknowable, or at least unknown—of our competent applications of 'rich' over the entire history of its use, and of course we cannot survey that entire history. (The sharp boundary is unstable insofar as the usage of a vague word may change over time.) In other words, the epistemicist contends that 'rich' has unknowable sharp boundaries that are fixed by an unknown function of the unknowable history of its competent use. Our use of the predicate is successful because competence does not require knowledge of the boundary's location:Competent speakers are permitted to make (what are strictly speaking) errors in their classifications of borderline cases. A contextualist variant of epistemicism holds that vague words have unknowable sharp boundaries whose locations shift from context to context (Graff 2000). On this view the major premise of the paradox is false but it seems true because the sharp boundaries shift in such a way that they are never located where we are currently looking or judging. A claimed virtue of epistemic theories is their retention of bivalence as well as classical logic: Any salary either is rich or is not, and the sentence '$*n* is rich' either is true or is false, though we cannot always know which.

The preceding accounts have advanced our understanding of vagueness, to be sure, but each is problematic in one way or another.[10]

For instance, the nonstandard conceptions of validity employed by multivalued theories, and indeed the very idea that truth comes in degrees, may seem unintuitive and/or ad hoc. The three-valued approach described above yields no tautologies. Also, a degree theory seems unmotivated in the case of a nongradable vague word like 'medium' (more on gradability shortly). Supervaluationism too has its shortcomings. In preserving excluded middle, the supervaluationist must say that a disjunction can be definitely true even if neither disjunct is definitely true; and it turns out that, when applied to sentences containing the definiteness operator, a number of classical inference rules including contraposition, conditional proof, and *reductio ad absurdum* no longer hold. Contextualist theories have come under fire for employing an ad hoc and unintuitive notion of a context, for applying only to a subset of vague words,[11] and for being improperly psychologistic. And the faith-based epistemic view multiplies mysteries.

Proponents of these theories have responded to many of the criticisms brought against them, and improvements have been made. But it remains the case that no theory of vagueness has yet won widespread acceptance. Part of the reason, I think, is that philosophers working on vagueness have taken for granted that a semantic theory—that is, a theory that regards vagueness as a semantic feature of language—cannot be classical.[12] In particular, they suppose that only an epistemic theory can preserve bivalence. Dominic Hyde writes that

> even if an epistemic analysis is possible, the indeterminacy surrounding the application of soritical terms is generally considered to be a semantic phenomenon. . . . If seen in this way, classical semantics appears in need of revision, and with it classical logic. (2011)

Williamson says,

> Most work on vagueness has taken it for granted that [epistemicism is] absurd. It therefore rejects the . . . supposition that an utterance of 'TW is thin' [where TW is a borderline case] is either true or false. (1994, 185)

The assumption is that bivalence entails sharp boundaries. Sven Rosenkranz asserts that "if all borderline statements are either true or false, as epistemicists contend, then vague words have sharp boundaries" (2003, 449). According to Rosanna Keefe, "epistemic theorists retain a classical [semantics]. This commits them to sharp boundaries to the extensions of our predicates" (2000, 62).

The theory to be developed here overturns this assumption. It is semantic—vague words lack sharp boundaries of application—but classical. (Contra Williamson, one can find epistemicism absurd without rejecting bivalence.) It is also simpler than its competitors in several respects; for example, it has no need of a definiteness operator, and while it recognizes the existence of higher-order vagueness, it rules out the possibility of higher-order borderline cases. This semantic but classical approach comes into view when we ground our theorizing about vagueness as deeply as possible in commonsense intuition and our actual competent use of vague words. Insofar as the sorites paradox can be generated from simple, intuitive propositions and rules of inference, we may expect intuitive considerations to play a central role in its solution. I think commonsense is rich with insights about vague words and that some of these have been overlooked. I will retrieve as many of them as I can and put them to theoretical use.

First I want to make some observations about vague words that should be more or less neutral among the theories described above. Some of these points will be familiar; some will not.

1.2. INITIAL OBSERVATION 1: BLURRED BOUNDARIES, SHARP BOUNDARIES, AND STOPPING PLACES

The distinction between blurred and sharp boundaries of application has been understood in a variety of ways.[13] As we saw a moment ago, many theorists cast it in terms of bivalence. Gary Ebbs writes,

> A concept has sharp boundaries . . . just in case for all *x*, the term that expresses the concept is true of *x* or false of *x*. The law of excluded middle entails that for all *x*, *x* is bald or *x* is not bald; if all sentences in which the word 'bald' occurs are bivalent, then the concept bald has sharp boundaries. (2001, 306)[14]

Wright characterizes blurred boundaries in terms of soriticality:

> Lack of sharp boundaries *as such* . . . seem[s] to imply paradox. To say that F lacks sharp boundaries in a series of the germane kind is to say, it seems, that there is no element x, which is F but whose immediate successor, x', is not. That is a claim of the form,
>
> $$\sim(\exists x)(Fx \,\&\, \sim Fx')$$
>
> and is accordingly classically equivalent to the major premise [of the sorites paradox,]
>
> $$(x)\,(Fx \rightarrow Fx').$$ (2007, 4)

The blurred/sharp distinction has also been analyzed in terms of a word's possession of borderline cases. Keefe writes,

> Clearly, having fuzzy boundaries is closely related to having borderline cases. More specifically, it is the *possibility* of borderline cases that counts for vagueness and fuzzy boundaries, for if all

> actually borderline tall people were destroyed, 'tall' would still lack sharp boundaries. It might be argued that for there to be no sharp boundary between the *F*s and the not-*F*s just *is* for there to be a region of possible borderline cases of *F* (sometimes known as the penumbra). (2000, 7)

Michael Tye seems to think that a sharp boundary would be a unique, fixed division that enjoys consensus among competent speakers. He writes,

> It seems clear that competent language users will not agree upon precisely where the boundaries are to be drawn in the sequence between the true, the indefinite, and the false statements. Of course, this is not to say that such people will not specify precise points if they are *forced* to assign either 'true' or 'false' or 'neither' to each of the statements . . . one after another. Still it seems highly unlikely that even one and the same person will pick exactly the same points on different occasions. It is not true, then, that the transitions from true to indefinite statements and from indefinite to false statements are sharp. (1994, 199)

And of course, epistemicists think that sharp boundaries are (unstable) divisions that are hidden from competent speakers.

However sharp boundaries are understood, they must be distinguished from mere permissible stopping places in a sorites series. By 'permissible stopping place' I mean simply any place at which a competent speaker, classifying the items in a sorites series *seriatim*, could permissibly stop applying the predicate in question. For example, in our series of salaries from $200,000 (clearly rich) to $50,000 (clearly not-rich), a competent user of 'rich' may permissibly stop applying it at $145,999, or at $145,998, or at $140,000, or at $136,002, among many others, even relative to the single context of Americans aged

forty to sixty in 2001.[15] All of these salaries and many others are permissible stopping places. As far as I know, no theorist of vagueness, including the epistemicist, thinks that the existence of permissible stopping places in a sorites series indicates the presence of sharp boundaries—not even given that stopping somewhere is a requirement of competence with the predicate.[16]

Perhaps if we identify some distinctive features of permissible stopping places we can understand sharp boundaries by contrast. The following seem like plausible candidates:

(i) Competent application of a vague word requires that users stop applying it before the end of a sorites series. Hence there must be some permissible stopping places.

(ii) Permissible stopping places are evident in the verbal behavior of competent users of a vague word. A competent user proceeding along a sorites series will in fact stop applying the word before reaching the last item.

(iii) However, no particular stopping place is required. Competent speakers are permitted to vary in their stopping places even relative to a single context. Different speakers, and the same speaker on different occasions, will stop at different places, even when contextual factors are held constant. (In chapter 5 I will provide some experimental confirmation of this variation; grant it for now.)

(iv) Accordingly, permissible stopping places are not legislative. Stopping at a particular place does not signify that other stopping places are incompetent or incorrect or even legitimately questionable. Multiple stopping places are equally permissible.

(v) Accordingly, any particular stopping place is arbitrary: One could always have stopped elsewhere. (Even the epistemicist will agree with this much.)

To say that a particular stopping place is arbitrary is not to say that it is irrational or unintelligible. Rationality does not always demand reasons. I will return to this point.

Our thought was that we might understand sharp boundaries by contrast with (mere) permissible stopping places. We will learn more about both notions as our theory of vagueness develops, but we can say straight off that sharp boundaries are not evident in the competent use of a vague word. Also, sharp boundaries (if such there be) are presumably legislative: If a speaker proceeding along a sorites series stops at a place other than the sharp boundary between the extension and anti-extension of the predicate, her stopping place is incorrect, strictly speaking.[17] The difference between sharp boundaries and mere stopping places will play a leading role in the theory of vagueness.

1.3. INITIAL OBSERVATION 2: VAGUENESS AND GRADABILITY

As we have seen, vagueness has been associated more or less closely with tolerance, borderline cases, soriticality, indeterminacy or indefiniteness, and ignorance, among other things, and I will say more about all of these. Vagueness has also been associated with gradability, about which I want to make a few early remarks.

In general a term is gradable just in case it permits comparisons. It does this by taking comparative and superlative forms and/or modifiers like 'more' and 'less' and/or intensifiers like 'very,' 'so,' and 'such.' The predicate 'tall' is gradable: One tree can be taller than another, and a building can be the tallest in the world; a person can be very tall, or too tall to play the leading role, or insufficiently tall to reach the cookie jar. In contrast, a person cannot be very 6 feet tall, or so next in line, or too dead. Kennedy explains

that "gradable adjectives establish relations between objects and measures of the degree to which they possess some property" (2004, 1). One thing can "possess tallness" to a greater or lesser degree than another. Arguably, nouns and adverbs also can be gradable: Stravinsky was such a genius, and Federer plays tennis more skillfully than golf.

Some theorists take gradability to be characteristic or even constitutive of vagueness. For example, Kees van Deemter asserts that "vague descriptions are referring expressions that contain gradable adjectives" (2006, 1; see also Sassoon 2007 for discussion). Not all vague words are gradable, however.[18] For example, suppose that the height of the flame on a gas stovetop is controlled by turning a continuously adjustable knob. And suppose that three settings are marked around the knob—'low,' 'medium,' 'high.' Presumably the boundaries between the low flames and the medium, and between the medium flames and the high, are unclear in the way that indicates vagueness; the predicates 'low,' 'medium,' and 'high' are vague. 'High' and 'low' are also gradable: For any two nonidentical flames, one is higher and the other lower. There are highest and lowest flames, and a flame can be so high that it burns the sauce or too low to melt the butter. 'Medium,' on the other hand, doesn't seem gradable: We would not say, of two flames, that one is *more or less medium* than the other—closer to or farther from being medium, perhaps, but not more or less medium. Talk of a 'very medium' or 'insufficiently medium' or 'too medium' flame doesn't sound good either, unless you mean that the flame is very or insufficiently or too *close* to medium.[19]

The precise predicate '6 feet tall' helps to illustrate the difference between being *Φ-er* or *more Φ*, on the one hand, and being *closer to* Φ, on the other. In an obvious sense, a height of 5 feet is closer to a height of 6 feet than a height of 4 feet is, but 5 feet is not *more 6 feet* than 4 feet is. By the same token, it makes no sense to say that

someone is very 6 feet tall, or not 6 feet tall enough to reach the cookie jar. Thus, although both the vague predicate 'medium' and the precise predicate '6 feet tall' are associated with a linear ordering of values on a decisive dimension of application, neither predicate is gradable. Mere affiliation with such an ordering does not guarantee either gradability or vagueness. (Unlike 'medium', the predicate '6 feet tall' is precise because there is no relevant unclarity about where its application begins and ends.)

Notice that despite being nongradable, 'medium' does seem to be soritical: If a given flame is medium, then another flame incrementally higher or lower is also medium. So gradability is not necessary for either vagueness or soriticality. The question of whether gradability is sufficient for either vagueness or soriticality is harder to answer; I will come back to it in chapter 4.

1.4. INITIAL OBSERVATION 3: VAGUENESS AND SORITICALITY

Vagueness is often defined in terms of soriticality, but this popular view doesn't stand up well under pressure. Most theorists of vagueness believe that the sorites paradox is a resolvable fallacy: The argument is defective, and we can figure out why. If that's right, then vagueness is not, after all, a source of paradox. Maybe someone will say that even after the defect in the sorites is discovered, the argument will still *appear* to consist in unimpeachable reasoning from true premises to a false conclusion and so will still be paradoxical.[20] But this is a strange view, for it makes vagueness a contingent psychological feature of language, dependent on how things appear to us. Furthermore—here I can judge only from my own case—once I had discovered (what I will argue is) the key to solving the puzzle, the major premise of the argument no longer appeared—no longer

appears—true. I can now still see why the premise formerly seemed true, or why it may seem true to the uninitiated, but again, this is a dubious way to define vagueness—namely, as the property of generating an argument that previously appeared, or appears to the uninitiated, paradoxical.

Perhaps the thought is that vagueness consists in generating an argument of a certain form.[21] Maybe this is what Bueno and Colyvan intend when they say that "a predicate is vague just in case it can be employed to generate a sorites argument" (2012, 29). But what does 'can be employed' mean here? If a sorites argument is a fallacy, a vague predicate cannot be *correctly* employed in it. Is the criterion supposed to be that a vague predicate is a term that, when employed incorrectly, appears (to the uninitiated) to generate an argument of the relevant form?[22] Surely vagueness is a more substantial property than that. Soriticality is an illusory feature of words like 'tall' and 'rich'; their vagueness is real.

1.5. INITIAL OBSERVATION 4: VAGUENESS AND CONTEXT-SENSITIVITY

Whatever the prospects for contextualist theories of vagueness, there is no denying the fact that most if not all vague words are contextually sensitive. A tree may be tall compared to aspens but not compared to redwoods; a country may be culturally rich but not financially so; a person's hair may be red in contrast to (as opposed to) brown but not in contrast to orange. If I tell you that someone or something is tall or rich or red, and you are not otherwise aware of the context in which I am speaking, you may not yet have enough information to assess the truth of my assertion. Hans Kamp writes,

> It is typical of a vague predicate that what objects it is true of depends on the context in which it is used. There are certain

> adjectives in particular—such as e.g. large, or soft, or clever—about whose extensions we can say hardly anything in abstraction from any contextual setting; it is only with respect to a given context of use that we can meaningfully ask whether a certain object is large, and there are very few, if any, objects of which it is clear absolutely whether or not they are clever or soft. Different contexts resolve these questions in different ways; the same object may count as definitely clever in one context and yet as definitely not in another. (1981, 242)

Context-sensitivity is probably not sufficient for vagueness: The expressions 'the fastest speed' and 'the winning score' seem context-sensitive but not vague (cf. Robertson 2000). Whether context-sensitivity is necessary for vagueness is a harder question; I will return to it later. However that issue is resolved, vagueness and context-sensitivity are frequent if not constant companions, so we need to understand how they are related. This question will be the focus of chapter 3.

1.6. VAGUENESS AND RULE-FOLLOWING

The preceding are just some initial, more or less theory-neutral observations about vague words. It will be helpful to have them on the table as we construct our theory. A principal aim of this book is to sift carefully through the phenomena of vagueness, soriticality, borderlines, context-sensitivity, and the rest to determine how, if at all, they are related. We will find that although they often appear together, they are less closely connected, and sometimes less prevalent in natural language, than is commonly supposed.[23]

We will also learn that a certain traditional conception of linguistic competence is incorrect. This traditional view—Wright calls it the

governing view (1987)—can be expressed as a pair of theses. The first is that language mastery is a fully rule-governed competence, and knowledge of the rules in question is implicit propositional knowledge (as opposed to some sort of mere "know how"). As Wright puts it, "linguistic competence is constituted by sensitivity to the dictates of internalized rules" (1987, 210). The second thesis is that the rules in question are discoverable by rational reflection, independently of any appeal to "external behavioural notions." Specifically, we can discover these rules by reflecting upon

> our conception of what *justifies* the application of a particular expression[,] our conception of what we should count as an adequate explanation of the sense of a particular expression[, and] the limitations imposed by our senses and memories on the kind of instruction which we can actually implement, [among other things]. (1975, 326–327)

Wright contends that the governing view is incoherent because, in the case of vague predicates, the rules served up by the reflective procedures it sanctions are inconsistent: One rule prescribes tolerant application of 'Φ,' while another prescribes application of 'not-Φ' to at least some items in a sorites series.[24] Together these rules generate the sorites paradox. Since our ordinary use of vague predicates is largely successful, Wright reasons, it cannot be guided by inconsistent rules (1975, 329). Therefore, if the use of vague words is in fact governed by rules, they are not rules of which we have implicit propositional knowledge or that we can discover by self-reflection. Wright resigns himself to the conclusion that mastery of a vague word is better understood "on the model of a practical skill, comparable to the ability to hit a good cross-court backhand or ride a bicycle" (1987, 239). In other words, mastery of a vague word is better understood as a kind of *knowledge how* than as *knowledge that*.

Like Wright, I will argue that the governing view is flawed, but the flaws I see are different from, and less damaging than, the ones he cites. Contrary to the first thesis of the governing view, linguistic competence is not entirely governed by rules; at a certain point, the rules give out and competent linguistic practice must become arbitrary. That fact notwithstanding, most of our use of vague words *is* rule-governed, hence not arbitrary; the rules in question are consistent, and our knowledge of them can be understood as implicit propositional knowledge. Moreover, these rules are discoverable largely by the self-reflective procedures described in the second thesis. As we will see, what's to blame for the sorites paradox is not the governing view so much as an overly strong conception of tolerance.

1.7. TWO POLICIES AND A CAUTION

I adopt two policies in what follows. First, probably many parts of speech are vague, including verbs (e.g., 'run'), adverbs (e.g., 'quickly'), and quantifiers (e.g., 'most');[25] but predicates—adjectives and nouns that predicate properties of things, as in 'The patch is blue' and 'That collection of sand grains is a heap'—have been of greatest interest to philosophers, and I will talk mostly about those. (In deliberately sloppy fashion I will speak indifferently about vague predicates and vague words.) In fact, much of the time I will work just with some lexical, noncomparative adjectives and nouns—familiar examples like 'rich,' 'old,' 'blue,' and 'heap.' I believe that the theory developed in this book is applicable to all vague terms, but I will not try to show that in a systematic way here.

Second, thus far I have used the term 'tolerance' in the usual way, following Wright, to refer to a semantic property of applying across incremental change on a decisive dimension. However, I will not use the term again until the end of the book, late in chapter 5; until then

I will speak only of the soriticality of vague words (keeping in mind that, ultimately, soriticality is illusory). This is because I think that vague words are tolerant, but the usual way of understanding tolerance is mistaken; in particular, tolerance is not properly expressed by the major premise of the sorites paradox. In chapter 5 I will say what I think tolerance consists in. It turns out to be a feature of the competent use of a vague word, as distinct from the semantics strictly speaking.[26]

Lastly a caveat: It is not the job of a theory of vagueness to deliver a verdict as to the vagueness, precision, soriticality, and so forth of every word that comes along. Certain questions may need to be put to competent linguistic intuition before any theory can be applied. For example, is the word 'strawberry' vague? Depending on one's theory, answering this question may require first determining whether 'strawberry' has borderline cases, or whether competent speakers are permitted to vary arbitrarily in its application, or whether there can be two incrementally different plants such that one is a strawberry and the other is not. As far as I can see, the latter questions need not, and often cannot, be answered by a theory of vagueness. Rather, they must be answered by competent, perhaps even expert, linguistic intuition.[27] (Maybe the botanists will have to be consulted.) We will see examples of this.

1.8. SELECTIVE REVIEW

Because my approach to vagueness departs substantially from the well-known treatments outlined above, and because coming to grips with an unfamiliar approach to a familiar topic is never easy, I will close each chapter with a review of certain key points. Such a review is needed least in this first chapter, so I will be brief. Let me reemphasize the following points:

(1) The widespread assumption that a semantic theory of vagueness cannot preserve bivalence is incorrect. The theory to be developed here is both semantic and bivalent. (It is *not* contextualist.)
(2) If the sorites paradox is solvable, a definition of vagueness in terms of soriticality cannot be correct.
(3) There are multiple equally permissible stopping places in a sorites series. (Chapter 5 will provide experimental evidence to support this claim.)
(4) Because there are multiple permissible stopping places, any particular stopping place is arbitrary and nonlegislative. One could always have stopped elsewhere.
(5) Some vague words, like 'medium' for example, are not gradable. Hence vagueness cannot be defined in terms of gradability.

1.9. LOOKING AHEAD

I said that an adequate theory of vagueness must resolve the sorites paradox, supply proper analyses of 'borderline case' and of 'vague' itself, and also show how vagueness, borderlines, and soriticality are related. Because my analysis of 'borderline case' introduces some machinery that I will need to use throughout the rest of the book, I will begin there.

Chapter 5

The Competent Use of Vague Words

> *Vague.*
> ORIGIN Latin *vagus* 'wandering'.
>
> *The Oxford English Dictionary*

5.1. A PRAGMATIC SORITES

Following Terence Horgan (e.g., 1994) and Scott Soames (1998), we can distinguish a secondary version of the sorites paradox, the so-called *forced march* or *dynamic* sorites. The dynamic paradox is an informal version of the argument framed in terms of the hypothetical applications of a vague predicate by a competent speaker who proceeds step by step along a sorites series. Consider a blue/green sorites series of thirty incrementally different patches (hues) and a competent speaker who begins at patch #1 and classifies each successive patch as blue, or as green, or as blue[green] borderline. On pain of incompetence, she must classify patch #1 as blue. Then since #2 is so similar to #1, it seems she must classify #2 as blue; and then #3 as blue, and so on until finally she must classify #30 as blue. The dynamic version of the paradox can be expressed this way:

> For any competent speaker *S* proceeding along the sorites series:
> (1D) *S* must classify patch #1 as blue.
> (2D) For any *n*, if *S* classifies patch #*n* as blue, then *S* must also classify patch #(*n* + 1) as blue.
> (3D) Therefore, *S* must classify patch #30 as blue.

At the same time, and also because she is competent, our speaker is bound to stop applying 'blue' before she reaches patch #30. Thus she is torn between apparently conflicting rules: She must violate either the rule expressed by the major premise (2D) or the rule that says that objects that look like patch #30 are green. Of course she will stop applying 'blue' before the end, and rightly so, but how, intuitively, does she manage to do this without disturbing the seamlessness of the series? I suggested at the end of the preceding chapter that we might think of the dynamic argument as a pragmatic version of the paradox, that is, a version that arises from the way we *use* vague words; and a solution to the semantic version, which we think we have now got, is not yet a solution to the pragmatic one. Only by resolving the pragmatic sorites we will fully understand how speakers manage to shift categories in an apparently seamless series—and indeed why, if they are false, the major premises of both versions of the puzzle seem true.

We can begin to resolve the dynamic sorites by attending to three facts about the competent use of a vague word. (I will continue to use 'blue,' 'green,' and 'blue[green] borderline' as my examples, but they are just examples.)

(1) First, each of the categories *blue*, *green*, and *borderline* contains a range of more and less central cases. Again, this is not the distinction between clear and borderline cases; the distinction between more and less central cases is drawn within each category: There are more and less central blues, more and less central greens, and more and less central borderline cases. In particular, as our speaker moves along the series away from the initial, most central blue patch (#1), the blue patches start to look more and more like the borderline patches. The result is that some patches that she competently classifies as blue and some patches that she competently classifies as borderline look

very much alike—even the same. (Remember that any patch that can competently be classified as borderline can also competently be classified as blue and as green.) This fact helps to demystify the speaker's ability to stop applying 'blue' and shift to 'borderline' between hues that are so similar. I suspect that when we think about the difference between being blue and being borderline, we think about the difference between looking like patch #1, a central case of blue, and looking like, say, patch #15, a central borderline case. So naturally it seems incredible that two neighboring patches could be one blue and the other borderline (a fortiori one blue and the other green.)

(2) Appeal to this first fact does not by itself dissolve the mystery of the dynamic sorites, since the possibility remains that all competent speakers stop applying 'blue' at the same place on every run along the series. For example, maybe everyone always shifts from 'blue' to 'borderline' at patch #13. This would suggest the presence of a sharp boundary. Hence we need to take account of a second, now familiar fact about the competent use of a vague word, namely, that it varies arbitrarily both within and across competent speakers, even with respect to a single V-index. In particular, different speakers diverge, and individual speakers vary, in their stopping places in a sorites series. In the terms of our semantic theory, a vague predicate has multiple arbitrarily different ranges of application relative to any of its V-indices. This arbitrariness is what enables variation in our use of a vague word without question of error.

(3) Even the first two facts together do not fully resolve the dynamic paradox, it seems to me. To see why, suppose that on a given run our competent speaker shifts from 'blue'

to 'borderline' at patch #15. Then, instead of having her continue along the series to #30, we immediately query her again about patch #14. If she still classifies #14 as blue, and especially if she persists in classifying #15 as borderline and #14 as blue, she would seem to have drawn a sharp boundary (if an idiosyncratic and temporary one) between #14 and #15. However, I predict that she will not classify #14 as blue after shifting to 'borderline' at #15; rather, she will now classify #14 as borderline too. The thought is that when she shifts from 'blue' to 'borderline' at #15, the speaker will undergo a characteristic change in her verbal dispositions so that she is now disposed to judge patch #14 as borderline also, despite having judged it blue just a moment before.[1] If she now continues backward down the series toward patch #1, she will classify some of the preceding patches as borderline before eventually shifting back to 'blue'; and so on. This is the third fact—well, so far it is only a hypothesis—that goes to dissolving the dynamic sorites. Intuitively, when the speaker categorizes patch #15 as borderline, it's not as if #14 still seems blue; instead, it's as if a string of patches shift their category together, so that consecutive patches never seem category-different at the same time (similarly for the later switch from 'borderline' to 'green'). Again, by the time the speaker shifts to 'borderline,' she is already in the midst of variable patches that can competently be called 'blue' and competently be called 'borderline' and competently be called 'green,' relative to a single fixed V-index. So the idea that her classifications should be fluid in this way is not so surprising.

The pattern of judgments that I have just described in (3) may constitute a *hysteresis* effect. As a first approximation, hysteresis is a

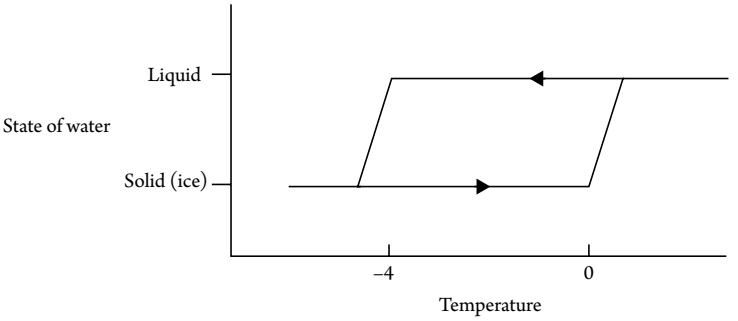

Figure 5.1. Hysteresis loop in the melting and freezing of water.

dynamical effect that occurs when an entity or system shifts discretely from one state to another as the result of an incremental change in the value of some determining parameter, and—here is the crucial bit—the value at the shifting point depends on the direction in which that parameter is changing. Van der Maas, Jansen et al. explain that hysteresis occurs "when the sudden jump [from one state to the other] depends on the direction of change in the normal variable. For instance, ice melts at 0°C but water freezes (in disturbance free conditions) at −4°C" (2004, 137). This pattern is commonly referred to as a 'hysteresis loop,' illustrated in Figure 5.1.

A classic case of hysteresis occurs in the behavior of magnets. The diagram in Figure 5.2 plots a hysteretic change in magnetization of a piece of iron. The *x*-axis specifies the strength of an externally applied magnetic field source (the 'determining parameter') and the *y*-axis specifies the degree of magnetization (the effect or 'response') in the iron. From the starting point, the magnetization of the iron increases as the applied magnetic field increases. But notice that when the magnetic field strength decreases from its maximal point, the magnetization of the iron does not simply retrace its steps back down to zero. Rather, as the field decreases, the iron retains its high level of magnetization for a period of time

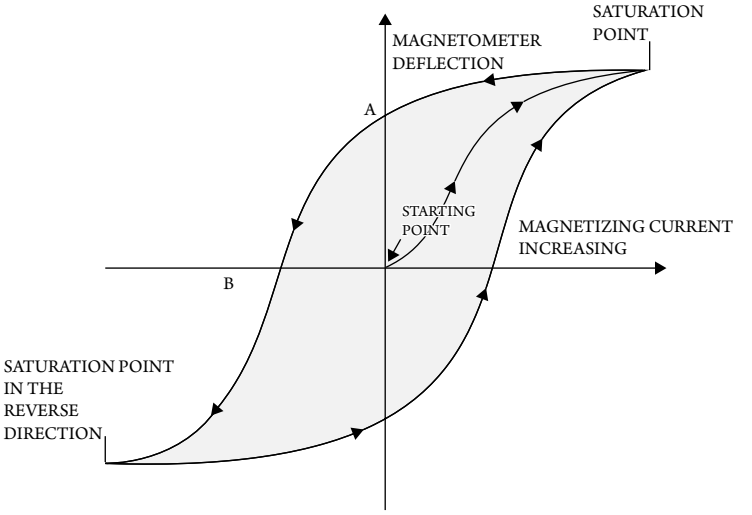

Figure 5.2. Magnetic hysteresis loop (http://www.daviddarling.info/encyclopedia/H/hysteresis_loop.html).

before it begins to drop. The decrease in magnetization lags behind the decrease in the field, and the magnetic field must be reversed and increased past zero in order to bring the iron back to its initial, non-magnetized state. (Hysteresis is sometimes called a 'lag' in, or 'persistence' or 'inertia' of, an effect or response.) Hence the iron's shift from a nonmagnetized state to a magnetized state occurs at a higher value of the applied magnetic field than its shift from magnetized to nonmagnetized.[2]

Hysteresis is also observed in human cognition, perception, and action. Just for example, Raczaszek et al. (1999) report a fascinating hysteresis effect in speakers' interpretations of ambiguous sentences as a function of continuously changing prosodic cues. (Prosodic features of an utterance include its intonation, rhythm, and stress pattern.) When disambiguating contextual information is absent, relative foot duration can disambiguate an utterance of a sentence

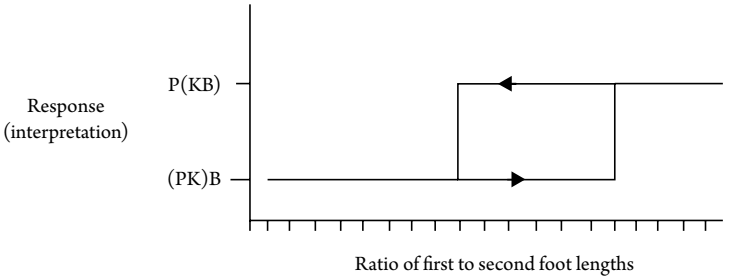

Figure 5.3. Hysteresis loop in interpretations of the sentence 'Pat and Kate or Bob will come.'

if the alternate interpretations of the sentence have different surface syntactic structures or "bracketings." Raczaszek et al. define a *foot* to be a 'string of syllables that begins with an accented syllable and extends to another accented syllable' (1999, 375). (So a foot is rather like a musical beat.) Consider for example the sentence 'Pat or Kate and Bob will come,' which has two possible interpretations and bracketings:

P(KB): [Pat] or [Kate and Bob] will come.
(PK) B: [Pat or Kate] and [Bob] will come.

The sentence has four feet: 1 = 'Pat or', 2 = 'Kate and,' 3 = 'Bob will,' 4 = 'come.' Raczaszek et al. observe that the ratio of the lengths of the first and second feet determines whether a hearer assigns interpretation P(KB) or interpretation (PK)B. In particular, the ratio of the first foot to the second is greater in P(KB) than in (PK)B.[3]

For their experiment Raczaszek et al. (1999) synthesized a series of utterances of the sentence 'Pat or Kate and Bob will come' that progressed by equal increments of foot duration from one that was heard as having the PK(B) interpretation to one that was heard as having P(KB). They found that subjects shifted from one interpretation to the other at different points in the series depending on the

direction in which the stimuli were changing. Specifically, subjects shifted from PK(B) to P(KB) closer to the P(KB) end than where they shifted from P(KB) to PK(B), so that their interpretations in the two directions overlapped (Figure 5.3). This is hysteresis.

Hysteresis is pervasive in the natural world. It is observed in ecological relations among predator and prey (e.g., Côté et al. 2004), in the functioning of labor markets and unemployment rates (e.g., Ball 2008), in decisions to buy or sell stock (Dixit 1992), in dating behavior (e.g., Tesser and Achee 1994), in heart rates (Walker et al. 2003), and in the process of cell differentiation (Kim 2008), to name just a few. Hysteresis is itself a certain pattern of behavior in a system, and instances of it in different systems result from different underlying mechanisms. However, a plausible general explanation for hysteretic behavior is that, all things being equal, many systems like to stay in the state they are currently in. (Among other things, changing from one state to another may expend energy.) I will discuss several forms of psychological hysteresis later. For now, I propose that once our speaker has shifted from 'blue' to 'borderline' in the blue/green dynamic sorites series, her categorizations will exhibit hysteresis: If she then reverses direction, she will categorize some preceding patches as borderline rather than blue as she had done before. The behavior of classifying patches as borderline will *persist* after, or *lag* behind, her category shift and change of direction (Figure 5.4). The idea is that it's *by* changing her dispositions in this hysteretic fashion that the speaker is able to shift categories while preserving the effective continuity, the seamlessness, of the sorites series.

You might think of these category shifts on the model of an automatic transmission in a car. Suppose that as the car accelerates from a stop, the transmission shifts from first gear to second at 20 mph. Once it has shifted to second, it will continue to use that gear as long as it can, even if the car subsequently slows to a speed previously handled by first gear: For example, the transmission will now remain

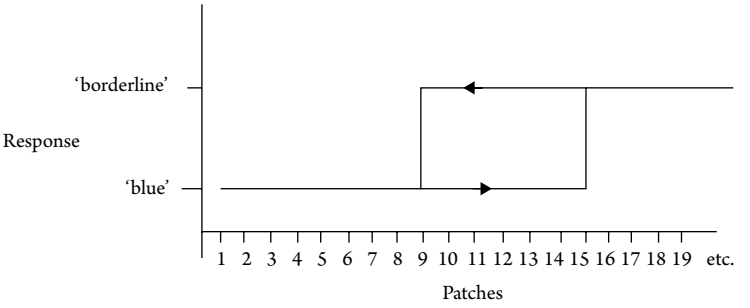

Figure 5.4. Predicted hysteresis in a blue/green dynamic sorites series.

in second gear even if the car slows to 15 mph, a speed previously handled by first gear.[4] The following explanation from a website on kart racing is helpful here:

> What is hysteresis? In racer's terms: (his-ta-'ree-sis) is the difference in engine speeds when the clutch engages and when it disengages. For example, an average clutch on an accelerating engine will engage at 4000 rpm. When slowing down from high speed the same clutch will disengage at 3500 rpm. This clutch has 500 rpm of hysteresis (4000 rpm–3500 rpm).... Hysteresis is a phenomenon in which the response of a physical system to an external influence depends not only on the present magnitude of that influence but also on the previous history of the system. Expressed mathematically, the response to the external influence is a doubled-valued function; one value applies when the influence is increasing, the other applies when the influence is decreasing.[5]

Analogously, once the competent speaker has shifted to 'borderline' and reversed direction, she will persist in using that category even when classifying patches that she formerly called 'blue.'

I don't mean to suggest that changing her verbal dispositions in this manner is under the speaker's conscious or otherwise willful control. For example, it's not as if she arrives at a certain patch and then *decides* to change her dispositions. (Maybe some subpersonal part of her "decides," but that is not something that she does or that she intends.[6]) Rather, when the speaker arrives at patch #15, it just strikes her a certain way—as borderline, for example. And if she reverses direction, some preceding patches will also now strike her as borderline even though they struck her as blue a moment before. The precise location of her category shift, and the associated change of verbal dispositions, must be the work of subpersonal mechanisms. The idea that speakers' applications of a vague predicate owe to the underlying switching machinery in this way accords with the arbitrariness of our stopping places in a sorites series; indeed, it may explain how that arbitrariness is possible. It also sits well with the absence of any normative framework for justification or argument concerning our classifications of variable items. I will say more about the role of mechanism in a theory of vagueness later.

Of course, to do the work I am assigning to it, the hypothesis of hysteresis will need to hold generally: It will need to hold for the competent use of any vague word, not just for predicates, and certainly not just for perceptual predicates like 'blue' and 'green.' For example, a speaker doing a forced march along a sorites series for 'rich' or 'old' or 'rich in contrast to middle income' should likewise (typically) undergo a hysteretic change of verbal dispositions at a category shift. If you proceed along our rich/middle income series of salaries $200,000 to $50,000 until you shift from 'rich' to 'borderline' at, say, $130,000, you should now (typically) be disposed to classify $130,001 and probably some of the preceding salaries as borderline also, in spite of the fact that you previously classified them as rich. Thus when you shift at $130,000, it's not as if you are crossing

a decisive threshold of richness; on the contrary, it's as if a string of salaries shift their category together, so that $130,000 and $130,001 never seem category-different at the same time. (Again, the 'seeming' needn't be phenomenological; the salaries will strike you, impress you, incline you to judge them, as rich or as borderline.) Remember that salaries in this transitional region of the series can competently be classified as rich and as borderline (hence not-rich) and as middle income, and your classifications of them will vary from occasion to occasion.

N.B. I am not proposing that hysteretic application is either necessary or sufficient for vagueness. Hysteretic application may be necessary for competent *use* of a vague word, but I am not going to argue for even that much here. I mean to propose only that hysteretic application may explain two things: first, how competent speakers are able to shift categories without disturbing the seamless progression of a sorites series, and second, how our multi-range semantics of vagueness could be implemented in competent use.

The idea that hysteresis occurs in a dynamic sorites series in the manner I am suggesting is an empirical hypothesis that needs to be tested. With two psychologists, Delwin Lindsey[7] and Angela Brown[8], I have designed and run an experiment to do that.

5.2. TESTING FOR HYSTERESIS

In designing our study, we faced the challenge of finding stimuli that could ensure that subjects were unaware of the reversal of direction immediately after a category shift. If subjects recognized the reversal, any spontaneous tendency to produce hysteretic judgments might be superseded by the desire not to appear inconsistent. Experimental subjects hate to appear inconsistent. At the reversal of direction, they would likely dig in their heels and think, 'I see what you are

doing—you are trying to catch me contradicting myself. But I'm smarter than you think. I'm going to draw a line and stick to it.' For this reason we could not employ sorites series of dollar amounts, or chronological ages, or numbers of sand grains. In fact, the only kind of stimuli we could think of that would serve our purpose were perceptual stimuli varying along a single continuous dimension. It is well known that such stimuli differing only incrementally cannot be reidentified across time. For instance, if I show you a blue patch, you will not be able to recognize it, as distinct from a patch of an incrementally different shade of blue, when you see it again even a few seconds later. In general, we can perceive (discriminate) many more values than we can recognize or even learn to recognize.[9]

With these considerations in mind, Lindsey constructed a set of thirty-seven colored lights progressing from a central blue to a central green, so ordered that each light looked the same in hue as the next.[10] The lights were viewed in a darkened room on a high-resolution hue monitor. Subjects' task throughout the experiment was to classify a single stimulus by clicking on one of three boxes labeled 'B' for blue, 'G' for green, and '?,' respectively. Subjects were instructed to "use the question mark response when, for any reason whatsoever, [they were] not fully satisfied either with 'B' or with 'G.'" The experiment was self-paced, but subjects were told not to spend more than a few seconds on each stimulus. The subjects were nineteen graduate and undergraduate students, postdocs, faculty, and staff from the philosophy and psychology departments at Ohio State University, including four experts in vision research. All testified that, to their knowledge, their hue vision was normal.

The experiment had three parts, each corresponding to a different way of presenting the stimuli. In Part I the lights were presented singly ("one at a time"), as illustrated in Figure 5.5. The colored light changed location in its box randomly from trial to trial, and the three responses 'B,' 'G,' and '?' switched boxes randomly between

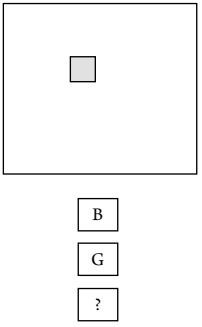

Figure 5.5. Stimulus configuration in Part I of the experiment. See www.oup.com/us/raffman for color illustrations.

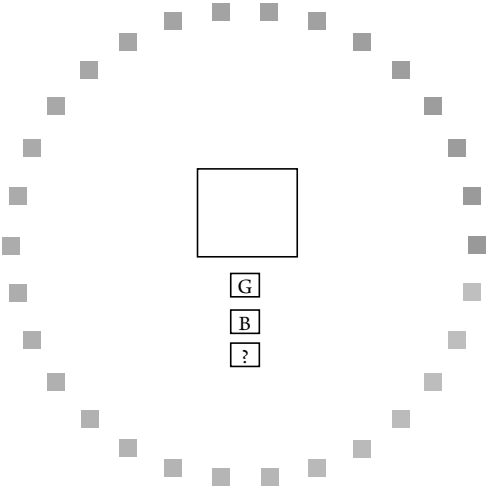

Figure 5.6. Stimulus configuration in Part II of the experiment. See www.oup.com/us/raffman for color illustrations.

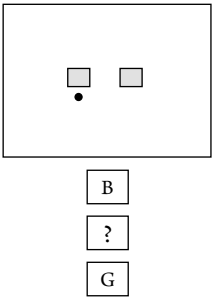

Figure 5.7. Stimulus configuration in Part III of the experiment. See www.oup.com/us/raffman for color illustrations.

conditions. Stimuli were presented in five conditions, each employing a different, randomly selected set of thirty consecutive stimuli from the available thirty-seven:

(1) Straight through the series in order from blue to green (BG)
(2) Straight through the series in order from green to blue (GB)
(3) Random order (Rn)
(4) Reversal ordering starting at the blue end (B[r])
(5) Reversal ordering starting at the green end (G[r]).

Each subject ran the set of five conditions twice. Within each set of five, conditions were presented in random order except that the first condition tested was always one of the two reversal orderings, selected at random. (We began each set with a reversal condition so as to minimize subjects' knowledge or memory of the stimulus set.)

Figures 5.8 and 5.9 (pp.150–151) display data from two individual subjects in Part I. Without defense beyond these few remarks, I have replaced the '?' response with lower-case 'b' for 'borderline.' We chose not to use 'borderline' as a response category in the experiment; since even the philosophers can't agree about what it means, we didn't want our subjects to be confused by it. Strictly speaking, then, the interpretation of the question mark remains open. The columns in each figure display the subject's responses in the five conditions, and each row represents the response on a given trial. The column labeled 'stm' specifies the stimulus being judged in a given trial, and the colored rectangles indicate the initial stimulus in each condition. The starting patch varied at random within a small range.

Notice first that our results provide evidence that competent speakers vary, both inter- and intrasubjectively, in their classifications of the items in a sorites series. For instance, in the first BG condition

stm	BG	GB	Rn	B[r]		G[r]			stm	BG	GB	Rn	B[r]		G[r]	
1	-	-	-	-	-	-	-		1	-	-	-	-	-	-	-
2	-	G	G	-	-	-	-		2	-	G	-	-	-	-	-
3	-	G	G	-	-	-	-		3	-	G	-	-	-	G	-
4	G	G	G	-	-	G	-		4	-	G	G	-	-	G	-
5	G	G	G			G	-		5	-	G	G	-	-	G	-
6	G	G	G	-	-	G	-		6	G	G	G	-	-	G	-
7	G	G	G	-	-	G	-		7	G	G	G	-	-	G	-
8	G	G	G	-	-	G	G		8	G	G	G	-	-	G	G
9	G	G	b	-	-	b	b		9	G	G	G	-	-	G	b
10	G	G	G	-	-	G	b		10	G	G	G	-	-	G	b
11	G	G	G	-	-	G	b		11	G	G	G	-	-	b	-
12	b	b	b	-	-	b	-		12	G	G	G	-	-	-	-
13	b	b	b	-	-	-	-		13	G	G	G	-	-	-	-
14	b	b	b	-	-	-	-		14	b	G	b	-	-	-	-
15	B	b	b	-	-	-	-		15	b	G	B	b	-	-	-
16	B	B	b	G	-	-	-		16	B	G	b	B	b	-	-
17	B	B	B	B	G	-	-		17	B	G	B	B	b	-	-
18	B	B	b	B	G	-	-		18	B	G	B	B	b	-	-
19	B	B	b	B	G	-	-		19	B	B	B	B	b	-	-
20	B	B	B	B	G	-	-		20	B	B	B	B	b	-	-
21	B	B	B	B	G	-	-		21	B	B	B	B	B	-	-
22	B	B	B	B	G	-	-		22	B	B	B	B	-	-	-
23	B	B	B	B	G	-	-		23	B	B	B	B	-	-	-
24	B	B	B	B	G	-	-		24	B	B	B	B	-	-	-
25	B	B	B	B	G	-	-		25	B	B	B	B	-	-	-
26	B	B	B	B	G	-	-		26	B	B	B	B	-	-	-
27	B	B	B	B	G	-	-		27	B	B	B	B	-	-	-
28	B	B	B	B	G	-	-		28	B	B	B	B	-	-	-
29	B	B	B	B	G	-	-		29	B	B	B	B	-	-	-
30	B	B	B	B	G	-	-		30	B	B	B	B	-	-	-
31	B	B	B	B	G	-	-		31	B	B	B	B	-	-	-
32	B	-	-	B	B	-	-		32	B	-	B	B	-	-	-
33	B	-	-	B	-	-	-		33	B	-	B	B	-	-	-
34	-	-	-	B	-	-	-		34	B	-	-	B	-	-	-
35	-	-	-	-	-	-	-		35	B	-	-	-	-	-	-
36	-	-	-	-	-	-	-		36	-	-	-	-	-	-	-
37	-	-	-	-	-	-	-		37	-	-	-	-	-	-	-

BG = stimuli presented in order (straight through) from blue to green
GB = stimuli presented in order (straight through) from green to blue
Rn = stimuli presented in random order
B[r] = reversal condition starting at blue
G[r] = reversal condition starting at green

Figure 5.8. Data from one subject for Part I of the experiment (single stimulus presentation). Curved arrows indicate reversal of direction and subsequent hysteresis. See www.oup.com/us/raffman for color illustrations.

in Figure 5.8, the subject shifts from 'blue' to 'borderline' at patch #14 and then from 'borderline' to 'green' at patch #11; whereas in the second BG condition he shifts from 'blue' to 'borderline' at #15 and from 'borderline' to 'green' at #13. In the first GB condition he shifts from 'green' to 'borderline' at #12, while in the second he shifts

stm	BG	GB	Rn	B[r]		G[r]			stm	BG	GB	Rn	B[r]		G[r]	
1	-	-	-	-	-	-	-		1	-	-	-	-	-	-	-
2	-	G	G	-	-	-	-		2	-	G	-	-	-	-	-
3	-	G	G	-	-	-	-		3	-	G	-	-	-	G	-
4	G	G	G	-	-	-	-		4	-	G	G	-	-	G	-
5	G	G	G	-	-	-	-		5	-	G	G	-	-	G	-
6	G	G	G	-	-	G	-		6	G	G	G	-	-	G	-
7	G	G	G	-	-	G	-		7	G	G	G	-	-	G	-
8	G	G	G	-	-	G	-		8	G	G	G	-	-	G	G
9	G	G	G	-	-	G	G		9	G	G	G	-	-	G	b
10	G	b	G	-	-	G	b		10	G	G	G	-	-	G	b
11	G	b	G	-	-	G	b		11	G	G	G	-	-	G	b
12	G	b	G	-	-	G	b		12	G	G	G	-	-	G	b
13	G	b	b	-	-	b	-		13	G	G	G	-	-	b	-
14	G	B	b	G	-	-	-		14	G	b	b	-	-	-	-
15	G	B	b	B	G	-	-		15	G	b	b	-	-	-	-
16	b	B	B	B	G	-	-		16	G	b	b	-	-	-	-
17	b	B	B	B	G	-	-		17	G	B	B	-	-	-	-
18	b	B	B	B	G	-	-		18	b	B	B	b	-	-	-
19	b	B	B	B	G	-	-		19	b	B	B	B	b	-	-
20	B	B	B	B	G	-	-		20	b	B	B	B	b	-	-
21	B	B	B	B	G	-	-		21	b	B	B	B	B	-	-
22	B	B	B	B	G	-	-		22	B	B	B	B	-	-	-
23	B	B	B	B	B	-	-		23	B	B	B	B	-	-	-
24	B	B	B	B	-	-	-		24	B	B	B	B	-	-	-
25	B	B	B	B	-	-	-		25	B	B	B	B	-	-	-
26	B	B	B	B	-	-	-		26	B	B	B	B	-	-	-
27	B	B	B	B	-	-	-		27	B	B	B	B	-	-	-
28	B	B	B	B	-	-	-		28	B	B	B	B	-	-	-
29	B	B	B	B	-	-	-		29	B	B	B	B	-	-	-
30	B	B	B	B	-	-	-		30	B	B	B	B	-	-	-
31	B	B	B	B	-	-	-		31	B	B	B	B	-	-	-
32	B	-	-	B	-	-	-		32	B	-	B	B	-	-	-
33	B	-	-	B	-	-	-		33	B	-	B	B	-	-	-
34	-	-	-	B	-	-	-		34	B	-	-	B	-	-	-
35	-	-	-	-	-	-	-		35	B	-	-	-	-	-	-
36	-	-	-	-	-	-	-		36	-	-	-	-	-	-	-
37	-	-	-	-	-	-	-		37	-	-	-	-	-	-	-

BG = stimuli presented in order (straight through) from blue to green
GB = stimuli presented in order (straight through) from green to blue
Rn = stimuli presented in random order
B[r] = reversal condition starting at blue
G[r] = reversal condition starting at green

Figure 5.9. Data from one subject for Part I of the experiment (single stimulus presentation). See www.oup.com/us/raffman for color illustrations.

directly from 'green' to 'blue' at #19. In Figure 5.9 we see still different classifications by another subject. On the plausible assumption that our subjects were competent speakers, these responses provide support for the claim that there are multiple equally competent stopping places in a sorites series—multiple equally competent ways of applying a vague word.

At present we are most interested in the reversal conditions, which tested for hysteresis. In reversal condition B[r], stimuli were presented starting from the blue end of the series. Immediately after the subject shifted to 'borderline' or to 'green,' the stimuli were then presented in reverse order. Thus if the subject shifted from 'blue' to 'green' at, say, stimulus #16, she was then presented with #17, #18, #19, and so forth until she shifted back to 'blue' (or 'borderline'). The G[r] condition did the same, starting from the green end. Both subjects display hysteresis. (See the curved arrows in the figures.) In the first reversal condition B[r] in Figure 5.8, for example, this subject classified lights #34 through #17 as 'blue' and then shifted to 'green' at #16. The next column then shows that as she was brought back down the series, she continued classifying lights #17 through #31 as 'green' even though she had judged them 'blue' just a moment before. In other words, she persisted in using the green category after her shift. We could say that in this block of trials she underwent a hysteresis of fifteen lights. In the second G[r] condition in Figure 5.9, we see that the subject shifted from 'green' to 'borderline' at light #13; he then displayed hysteresis of four lights until shifting back to 'green' at light #8. (Here is one way in which the idea that vagueness is a form of *wandering* may be embodied in competent use.[11])

Figures 5.8 and 5.9 show two of our most striking cases; often the hysteresis was shorter, around two or three lights. The average length of hysteresis observed in Part I of the experiment was slightly less than three lights. One expert did not exhibit hysteresis in any block of trials. There is nothing in my hypothesis to rule out that possibility—it is simply a limiting case of hysteresis of zero lights—but it should be, and was, exceptional.[12] Another striking aspect of these data is that subjects often did not use the '?' response (see, e.g., the first B[r] and second GB conditions in Figure 5.8). In fact our expert subjects almost never used the question mark. If '?' can indeed be

interpreted as meaning 'borderline,' this finding provides support for the idea, floated during our discussion of higher order borderline cases in chapter 2, that a judgment of 'borderline' is always optional; there are no definite borderline cases. Actually it indicates something stronger—namely, that a competent speaker needn't classify *any* items in a sorites series as borderline.

We ran the 'straight through' conditions BG and GB partly in order to see whether our study would replicate the results of a previous, standard test for hysteresis reported in Kalmus (1979). The standard way to test for hysteresis is to run subjects through an entire series of stimuli first in one direction and then in the other, as we did in BG and GB. Subjects' classifications are hysteretic if the shifting point in each direction is displaced toward the opposite endpoint—for example, if the span of patches classified as blue in the blue-to-green direction overlaps with the span of patches classified as green in the green-to-blue direction. Kalmus (1979) performed exactly this test with hue stimuli and found no hysteresis. Instead he found enhanced contrast, which is the opposite of hysteresis: Enhanced contrast consists in a gap, rather than an overlap, between the classifications in the two directions. Our study largely replicates Kalmus's finding. Notice for instance that in Figure 5.9, the last blue patch in the first BG condition is #20, whereas the last green patch in the first GB condition is #9. The last blue patch in the first B[r] condition is #15, whereas the last green patch in the first G[r] condition is #12. Thus, interestingly, we observed hysteresis only after a reversal of direction. A reversal condition has not been investigated previously, so far as we are aware.

In Part II, thirty consecutive stimuli were presented simultaneously, as shown in Figure 5.6 on page 148; all of the patches were in view throughout. On each trial, a small black dot appeared next to the patch that was to be classified. Data from one subject are shown

in Figure 5.10 below. These results bear out our expectation that if subjects were aware of the change of direction in the reversal condition, they would tend to select an arbitrary fixed boundary so as to avoid appearing inconsistent, and this would tend to obscure

stm	BG	GB	Rn	B[r]		G[r]			BG	CB	Rn	B[r]		G[r]		stm
1																1
2	G								G							2
3	G	G							G					G		3
4	G	G							G					G		4
5	G	G	G			G			G	G	G			G		5
6	G	G	G			G			G	G	G			G		6
7	G	G	G			G			G	G	G			G		7
8	G	G	G			G			G	G	G			G		8
9	G	G	G			G			G	G	G			G	G	9
10	G	G	G			G	G		G	G	G			b		10
11	G	G	G			b			G	G	G					11
12	b	b	b						G	G	G					12
13	b	b	b						G	G	b					13
14	B	B	b	b					G	b	b					14
15	B	B	G	B	B				G	b	b					15
16	B	B	B	B					b	B	b					16
17	B	B	B	B					B	B	B	G				17
18	B	B	B	B					B	B	B	B	B			18
19	B	B	B	B					B	B	B	B				19
20	B	B	B	B					B	B	B	B				20
21	B	B	B	B					B	B	B	B				21
22	B	B	B	B					B	B	B	B				22
23	B	B	B	B					B	B	B	B				23
24	B	B	B	B					B	B	B	B				24
25	B	B	B	B					B	B	B	B				25
26	B	B	B	B					B	B	B	B				26
27	B	B	B	B					B	B	B	B				27
28	B	B	B	B					B	B	B	B				28
29	B	B	B	B					B	B	B	B				29
30	B	B	B	B					B	B	B	B				30
31	B	B	B	B					B	B	B	B				31
32		B	B	B						B	B	B				32
33			B							B	B	B				33
34			B							B	B					34
35																35
36																36
37																37

BG = stimuli presented in order (straight through) from blue to green
GB = stimuli presented in order (straight through) from green to blue
Rn = stimuli presented in random order
B[r] = reversal condition starting at blue
G[r] = reversal condition starting at green

Figure 5.10. Data from one subject for Part II of the experiment (simultaneous presentation). No hysteresis evident. See www.oup.com/us/raffman for color illustrations.

any hysteresis in their classifications. Nevertheless, some subjects did exhibit a short hysteresis in Part II, as indicated in Figure 5.11, though the average length was only slightly more than one patch. As in Part I, the '?' response was not always used.

Perhaps it will be thought that subjects' judgments exhibited hysteresis in Part I merely because they believed or expected that the stimuli they saw in the reversal conditions B[r] and G[r] were continuing to change in the original direction. For instance, after shifting from 'blue' to 'green,' maybe subjects persisted in saying 'green' not because the stimuli looked green but only because the stimuli had so far been progressing steadily from blue toward green, and so subjects expected them to continue in that direction. After a time, as the stimuli became increasingly blue, subjects realized that the direction had reversed and went back to saying 'blue.'

This is an intelligent objection, and we cannot rule out the possibility it describes. However, we can offer some reasons to favor a hypothesis of perceptual hysteresis. First, in Part II of the experiment, all of the stimuli were presented simultaneously, so subjects presumably did not believe that stimuli were continuing to progress in the same direction. (Call the latter belief the 'continuation belief.') Nevertheless, six subjects displayed a hysteresis of 1 patch, and two displayed a hysteresis two patches long. Hence a continuation belief is unlikely to have been the cause of the hysteresis in either Part I or Part II. Also, in some instances in Part I, the hysteresis was long enough to make it unlikely that the subjects still believed the stimuli were progressing in the original direction, yet they persisted in the new (post-shift) category anyway. Three subjects exhibited hysteresis more than eight patches long, and the subject represented in Figure 5.8 exhibited a hysteresis of fifteen patches. These results suggest that the continuation belief does not explain the hysteretic patterns of judgment we observed.[13]

stm	BG	GB	Rn	B[r]		G[r]			stm	BG	GB	Rn	B[r]		G[r]	
1	-	-	-	-	-	-	-		1	-	-	-	-	-	-	-
2	-	-	G	-	-	-	-		2	G	G	-	-	-	-	-
3	-	-	G	-	-	-	-		3	G	G	-	-	-	-	-
4	G	-	G	-	-	-	-		4	G	G	G	-	-	-	-
5	G	G	G	-	-	-	-		5	G	G	G	-	-	G	-
6	G	G	G	-	-	G	-		6	G	G	G	-	-	G	-
7	G	G	G	-	-	G	-		7	G	G	G	-	-	G	-
8	G	G	G	-	-	G	-		8	G	G	G	-	-	G	-
9	G	G	G	-	-	G	-		9	G	G	G	-	-	G	-
10	G	G	G	-	-	G	-		10	G	G	G	-	-	G	-
11	G	G	G	-	-	G	-		11	G	G	G	-	-	G	-
12	G	G	G	-	-	G	-		12	G	G	G	-	-	G	-
13	G	G	b	-	-	G	-		13	G	G	G	-	-	G	G
14	G	b	b	-	-	G	-		14	G	b	G	-	-	b	-
15	G	b	b	-	-	G	G		15	b	b	G	-	-	-	-
16	b	b	b	-	-	b	-		16	b	b	b	-	-	-	-
17	b	b	b	-	-	-	-		17	b	b	G	-	-	-	-
18	b	b	b	b	-	-	-		18	b	b	G	-	-	-	-
19	b	B	B	B	b	-	-		19	b	B	b	-	-	-	-
20	b	B	B	B	B	-	-		20	B	B	b	b	-	-	-
21	B	B	B	B	-	-	-		21	B	B	b	B	b	-	-
22	B	B	B	B	-	-	-		22	B	B	B	B	b	-	-
23	B	B	B	B	-	-	-		23	B	B	B	B	B	-	-
24	B	B	B	B	-	-	-		24	B	B	B	B	-	-	-
25	B	B	B	B	-	-	-		25	B	B	B	B	-	-	-
26	B	B	B	B	-	-	-		26	B	B	B	B	-	-	-
27	B	B	B	B	-	-	-		27	B	B	B	B	-	-	-
28	B	B	B	B	-	-	-		28	B	B	B	B	-	-	-
29	B	B	B	B	-	-	-		29	B	B	B	B	-	-	-
30	B	B	B	B	-	-	-		30	B	B	B	B	-	-	-
31	B	B	B	B	-	-	-		31	B	B	B	B	-	-	-
32	B	B	-	-	-	-	-		32	-	-	B	B	-	-	-
33	B	B	-	-	-	-	-		33	-	-	B	B	-	-	-
34	-	B	-	-	-	-	-		34	-	-	-	B	-	-	-
35	-	-	-	-	-	-	-		35	-	-	-	-	-	-	-
36	-	-	-	-	-	-	-		36	-	-	-	-	-	-	-
37	-	-	-	-	-	-	-		37	-	-	-	-	-	-	-

BG = stimuli presented in order (straight through) from blue to green
GB = stimuli presented in order (straight through) from green to blue
Rn = stimuli presented in random order
B[r] = reversal condition starting at blue
G[r] = reversal condition starting at green

Figure 5.11. Data from one subject for Part II of the experiment (simultaneous presentation). Small hysteresis in B[r]. See www.oup.com/us/raffman for color illustrations.

5.3. NONPERCEPTUAL HYSTERESIS: DOES OUR HYPOTHESIS GENERALIZE?

Even if applications of hue predicates are hysteretic in the way I predicted, that doesn't yet show that the same holds for nonperceptual

vague words like 'rich,' 'old,' and 'person'; maybe our experimental results were elicited by specifically perceptual features of the stimuli. I do not know of any studies of hysteresis effects in nonperceptual categorization in an immediate reversal condition like the one we investigated.[14] However, there is ample evidence of nonperceptual hysteresis effects in other conditions, so I think we can be optimistic that our hypothesis will hold up outside the perceptual domain. Let me provide a few illustrations. (Since it is important to see exactly what these theorists say, I am going to quote several of them at some length.)

With regard to attitudinal change, Van der Maas, Jansen et al. explain:

> Hysteresis in . . . attitude . . . means that the informational value at which people change their attitude (the threshold) depends on the person's initial position and the direction of change in information. For instance, people with a positive attitude toward abortion change to a negative attitude only when information is strongly against abortion. People with a negative attitude change to a positive attitude only when information is strongly in favor of abortion. (2003, 131)

Such attitudes toward abortion might be reflected in hysteretic application of vague predicates like 'morally (im)permissible' or 'sinful' or maybe 'person'. Oliva et al. describe hysteresis in the relationship between service delivery and customer satisfaction in marketing:

> Once a sudden shift occurs, a return to the former [behavior] will not occur even if the independent variable values return to the levels they were [at] when the shift was made. There is a lag or hysteresis (inertia) in the process, which tends to keep behavior at its current level. For example, when a brand loyal customer for a high involvement product switches to another brand because of dissatisfaction with the first brand's performance or service,

> the customer is not likely to be easily switched back. . . . The reason is that the [customer's] current behavior . . . is dependent on recent past [performance], just as service adjustments provided to an unhappy customer may not reduce unhappiness, and the same adjustments made to a happy customer may not reduce happiness. (1992, 87)

We can perhaps imagine customers applying a vague predicate such as 'reliable detergent' or 'effective cough suppressant' in a hysteretic fashion.

Abraham Tesser and John Achee (1994) have developed a model in which dating behavior in the face of negative social pressure (e.g., dating someone of the "wrong" gender, race, or religion) exhibits hysteresis. Here their model is discussed by Jonathan Elster:

> As the disposition [to date the "wrong" person] continues to decrease in the face of strong social pressure, there will come a point when the person switches from engaging in the behavior to not engaging in it. In the [opposite] case, as the disposition continues to increase, there will come a point when the person switches from not engaging in the behavior to engaging in it. Moreover, the level of disposition at which the first switch occurs is lower than the level at which the second occurs. A person who has adopted an unpopular opinion will need to see a lot of the evidence for it fritter away before he gives it up, whereas an uncommitted person will need a lot of evidence for it before adopting it. Finally, a given combination of social pressure and disposition can lead to high as well as low engagement in the behavior, depending on where the person initially started up. (1976, 68)

Here the vague terms implicitly at issue might be something like 'datable' or 'mate-worthy.'

Economists also are familiar with hysteresis in human behavior. For instance, Avinash Dixit asks us to

> picture a particular path of the stochastic evolution of net revenues [*R*] of a company through time. . . . Suppose the initial *R* equals 1, and it starts to rise. . . . Finally it rises above 1.62, and the project is launched. Then the revenue starts to fall, and comes back all the way down to 1. But this does not justify abandonment. The driving force behind the investment decision, namely the currently observed revenue, has been restored to its initial level. But its meandering along the way has left its mark, namely an active project where there was none before.
>
> Similar effects have long been known in physics and other sciences . . . This phenomenon is called hysteresis, and by analogy the failure of investment decisions to reverse themselves when the underlying causes are fully reversed can be called economic hysteresis. . . . If our project's current revenue falls even more, it will eventually be abandoned. Then a subsequent rise back to 1 in revenues will not restore the project; there is hysteresis in the reverse direction, too. . . .Very large changes in *R* in the opposite direction are needed to reverse the effects of a temporary move in either direction. (1992, 121–122)

Here the vague words at issue might be something like 'good (bad) time to invest' or 'adequate revenue to justify launching the new project.'[15]

These examples are not conclusive. As I have said, I do not know how to prove that hysteresis of the sort observed in our experiment, namely, hysteresis in an immediate reversal condition, would occur in applications of nonperceptual vague predicates as well as perceptual ones. Van der Maas et al. don't describe a case in which information about abortion turns negative immediately after an agent has

shifted from a negative to a positive attitude toward abortion, and Dixit (1992) doesn't describe a case in which revenue decreases immediately after a project is launched. Plus there remains the difficulty of revealing hysteresis when subjects are aware of the reversal. Thus the best I can offer at present is a plausibility claim: The research cited above suggests that a variety of nonperceptual judgments are hysteretic in "straight-through" conditions, and from there it is not a great leap to the idea that the dynamics of nonperceptual judgments over incremental differences would be relevantly analogous to those observed in our experiment. (Indeed, given that we observed hysteresis only in the reversal conditions, and not in the straight-through conditions, the application of predicates like 'datable,' 'morally impermissible,' and 'reliable detergent,' which are likely to be hysteretic even in straight-through conditions, might exhibit greater hysteresis than the hue words in a reversal condition.) Let us suppose, then, at least for the sake of argument, that our hypothesis holds of vague words across the board.

Wright has raised the worry (in conversation) that competent speakers might exhibit hysteresis in their (perceptually based) applications of some precise predicates also. They might not apply an obviously precise predicate like 'six feet tall' hysteretically; indeed they might decline to apply that predicate at all, just on the basis of looking. But they might apply a less obviously precise predicate hysteretically—say, the predicate 'taller than six feet.' Speakers might apply the latter predicate hysteretically even though arguably it is neither soritical nor vague. If that's right, our finding of hysteresis shows nothing distinctive about vague predicates.

Let me offer two points in reply. First, notice that Wright's worry cannot be allayed by contending that the speakers in his example would implicitly be applying not the predicate 'taller than six feet' but rather 'looks taller than six feet,' for this contention misses the distinction between judging that something looks Φ, on the one

hand, and judging, *on the basis of looking*, that something *is* Φ, on the other. We should suppose that the speakers in the example are doing the latter: They are judging, on the basis of looking, whether a given object is taller than six feet. However, Wright's worry *can* be allayed by appeal to the further distinction, often drawn by philosophers of perception, between an object's looking Φ and an object's looking as if it is Φ (or looking to be Φ). For instance, a round object viewed from an angle looks elliptical but does not look as if it is elliptical. It looks as if it is round; round things look elliptical when viewed from an angle. My point then is that if speakers applied the predicate 'taller than six feet' hysteretically, on the basis of looking, they would in fact be judging not that the object in question is taller than six feet, or even that it looks taller than six feet, but rather that it looks *as if it is* taller than six feet, or *looks to be* taller than six feet. (It looks the way things that are taller than six feet look.) Because they would be classifying the object on the basis of looking, they would actually, implicitly, be applying the predicate 'looks as if it is taller than six feet'; when it is applied *on the basis of looking*, the predicate 'taller than six feet' is simply an abbreviation for 'looks as if it is taller than six feet.' And the latter predicate is vague. Hence hysteresis is just what we should expect. Second, as I said above, I am not claiming that hysteresis in the application of a term is either necessary or sufficient for vagueness; at most, hysteresis may be necessary for competent use of a vague term (as opposed to competent use of a precise term). I am claiming only that evidence exists to suggest that competent speakers apply vague words hysteretically and that the hysteresis may explain how they are able to shift categories in a dynamic sorites series without disturbing its apparent continuity.

My hypothesis then is that hysteresis in our applications of a vague predicate to the items in a dynamic sorites series has the effect of smoothing out what would otherwise be sharp or abrupt transitions from one category to another. (There is in fact a mathematical

technique called 'hysteresis smoothing' that is used to smooth pixels in digital images and to enable smooth starting in clocks and turntables, for example.) Intuitively, hysteresis in our applications of a vague word allows us to talk about the continuous world around us; it allows our "discrete" language to apply to a seamless environment. Generally speaking, items in the transitional region between vague predicates 'Φ' and 'Φ*' can competently be classified as Φ and as Φ* and as Φ[Φ*] borderline. So we can hardly cast aspersions on our classifications of those items for being hysteretic. If anything, the hysteresis provides a pattern or regularity to our variable judgments. It makes them more intelligible, not less: The permissible variability is hysteretic, not random.

A referee for Oxford University Press asks, "[If the subjects] are trying to be consistent (as, it seems, some are), why not take *that* to be due to semantic features of ['blue' and 'green']? That is, why chalk it up to artificial features of the testing environment?" The thought here, I take it, is either (a) that 'blue' and 'green' have unknowable, or at least unknown, sharp boundaries that our inconsistent use is (unbeknownst to us) aiming at or (b) that the semantics of these predicates is such that although they lack sharp boundaries, our use of them aspires or ought to aspire to such boundaries, that is, we should try to apply them *as if* they had sharp boundaries, and that is why we try to maintain consistency in our applications when the stimuli are presented simultaneously. I cannot see any good reason to endorse either (a) or (b) and many reasons to reject both; and even if such reasons exist, they will be outweighed by considerations favoring the explanation I have offered. Let me mention four of these.

First, my explanation is not ad hoc: Subjects' desire to be consistent in their responses has been observed across a wide range of experimental conditions (see Falk and Zimmermann 2013, for example). Moreover, knowledge of that desire led us to make a

correct prediction about the results of Part II of the experiment: We predicted little or no hysteresis in the reversal conditions. Second, fully competent speakers are incapable of applying these predicates consistently. Why would their semantics require that we try to apply them in a humanly impossible way? Third, if the experimental findings are correct, the inconsistency in subjects' applications of vague words is not anomalous: It is hysteretic, and the hysteresis is what enables smooth category shifts. The existence of this pattern, and its apparent purpose, suggest that we apply vague words hysteretically *by design*, as one might put it. Fourth and last, if *per impossibile* we could apply 'blue' and 'rich' and 'bald' consistently, then as Wright points out (cf. p. 3 above), those words would be effectively useless in communication. Why would their semantics dictate that we aspire to apply them in a way that makes them useless?

I said that we could resolve the dynamic version of the paradox by attending to three aspects of vague predicates and their competent use: (1) each category in a sorites series encompasses a range of more and less central cases; (2) there are multiple equally competent ways to apply a vague predicate—in particular, multiple equally competent places to stop applying it in a sorites series; and (3) competent applications of a vague predicate are hysteretic in the manner observed in our experiment (we are supposing). It seems to me that, taken together, these features of competent use demystify our ability to shift categories without introducing an abrupt transition into a seamless sorites series. In so doing they show how the major premise of the dynamic paradox can be false: It is not true that if a speaker has applied the predicate in question to one item in a sorites series, she must also apply it to the next.

If the major premise of the dynamic sorites is false, indeed necessarily false, why does it seem true? I will offer an answer to this pressing question shortly.

5.4. MEANING AND USE: IMPLEMENTING THE MULTIPLE RANGE SEMANTICS

In developing the multiple range theory, I have placed significant weight on the idea that the character of the competent use of vague words provides evidence of their semantic structure. In particular, the hypothesis that vague words have multiple arbitrarily different ranges of application in their semantics is motivated in part by the variability of their competent use. Our experiment has now revealed additional aspects of the competent use of vague words, in particular its hysteretic character, that would not otherwise have been apparent. Can the hysteresis observed in the experiment also be seen as implementing the multi-range semantic theory, and if so, how?

Think of the use of vague words on the model of an imaginary board game. You have a game piece, and the game board contains an array of twelve colored patches progressing from a blue one to a green one, as illustrated in the top panel in Figure 5.12. (I have used twelve patches merely for practical reasons; it helps to imagine a more populated series having, say, thirty hues as in our experiment.) Suppose that the operative V-index for 'blue' is ⟨hue; green; 0⟩, the V-index for 'green' is ⟨hue; blue; 0⟩, and the V-index for 'borderline' is ⟨hue; blue, green; 0⟩. The object of the game is to produce a competent classification of all twelve patches seriatim, that is, to produce a dynamic sorites series, by landing your game piece on various squares in the rows (*i*-*x*) on the board. Rows α and β at the top of the game board illustrate two possible ways of classifying the patches.

As an exercise, let's try producing the classifications in row α (Figure 5.12, bottom panel). You begin at patch #1 and classify one patch on each move by "using" the rows on the game board. For example, you can classify patch #1 as blue ('B') using any of rows *i* through *iii*; let's suppose you choose row *ii*, so your place your game piece on

α	B	B	B	B	B	B	B	b	G	G	G	G	
β	B	B	B	B	B	b	b	G	G	G	G	G	

	1	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	
i	B	B	B	B	B								***i***
ii	B	B	B	B	B	B							***ii***
iii	B	B	B	B	B	B	B						***iii***
iv				b	b	b	b						***iv***
v					b	b	b						***v***
vi						b	b	b	b				***vi***
vii					b	b	b	b					***vii***
viii						G	G	G	G	G	G	G	***viii***
ix					G	G	G	G	G	G	G	G	***ix***
x							G	G	G	G	G	G	***x***

i	B	B	(B)	B	B								***i***
ii	(B)	(B)	B	B	B	(B)							***ii***
iii	B	B	B	(B)	(B)	B	(B)						***iii***
iv				b	b	b	b						***iv***
v					b	b	b						***v***
vi						b	b	b	b				***vi***
vii					b	b	b	(b)					***vii***
viii						G	G	G	G	(G)	G	G	***viii***
ix					G	G	G	G	(G)	G	G	G	***ix***
x							G	G	G	G	(G)	(G)	***x***

Figure 5.12. Top panel shows the game board. Bottom panel shows the rows that could be used to produce the classifications in row α, with a particular sequence of moves circled. See www.oup.com/us/raffman for color illustrations.

the first square of that row (#1: *ii*). You can also classify patches #2 through #5 as blue by using any of rows *i* through *iii*; let's suppose that you move your piece to the second square of row *ii* (#2: *ii*), then to the third square of row *i* (#3: *i*), then to the fourth square of row *iii* (#4: *iii*), and then to the fifth square of row *iii* (#5: *iii*). You can classify #6 as blue by placing your game piece on either row *ii* or row *iii*, and classify #7 as blue using row *iii*. Patch #8 is classified as borderline ('b') using *vi* or *vii*, and patches #9-#12 can be classified as green ('G') using any of rows *viii* through *x*. On this particular game board, you can classify patch #6 as blue or as borderline or as green, but you can classify patch #8 only as borderline or as green and patch #1 only as blue. A judgment of 'borderline' is always optional; there are no patches having only 'b' rows below them.

Granted *Monopoly* has nothing to fear, but this toy example helps us to see how the multi-range semantic theory might be implemented in our competent use of a vague word. Using a vague word in a dynamic sorites series may be similar in certain respects to playing the game. Think of the rows on the board as different ranges of application of 'blue,' 'borderline,' and 'green': Rows *i* through *iii* as ranges of 'blue,' rows *iv* through *vii* as ranges of 'borderline,' and rows *viii* through *x* as ranges of 'green.' Then think of classifying a patch using a given row as analogous to classifying an item relative to a corresponding range of application. Where a player is permitted to classify a patch as, for example, green by using a 'green' row containing the hue of that patch, a speaker is permitted to classify an item as green by "using" a range of application of 'green' that contains the hue of that item. Where the player traverses the board by jumping among overlapping rows of different categories, the speaker traverses a sorites series by "jumping" among different overlapping ranges of the relevant predicates, much as one might bridge a crevasse by jumping from one to another of a series of overlapping planks of different lengths.

We can also understand the hysteresis in speakers' applications of vague words by analogy with possible moves on the game board. Suppose you want to produce the assignments shown in row β. You would classify patches #1 through #5 using any of rows *i–iii*, patches #6 and #7 relative to any of *iv–vii*, and patches #8–#12 relative to any of *viii–x*. Now suppose that, for whatever reason, you decide to reverse direction after classifying #8 as green. The overlaps among the rows would enable a hysteretic pattern in the reversal. You would trace a hysteretic path if you classified some preceding patches as green also; for example, if you classified patch #7 using row *viii*, patch #6 using row *ix*, #5 also using *ix*, and then shifted back to blue at #4. (See the bottom panel of Figure 5.13.) By analogy, a speaker could classify patch #7 as green relative to ("using") range *viii*, patch #6 using range *ix*, and so forth.

The analogy goes only so far, however. In particular, whereas the player in the game selects a single row to use on each move, it does not seem plausible that, in any given application of a vague word, a competent speaker uses or relativizes to a particular unique range of application. What could determine a particular range? When I say that the sky is blue, what facts about me or the language or the world could determine that I am applying 'blue' relative to, for example, range *i* rather than range *ii* or range *iii*? I think we should say instead that for any V-index, a speaker applies a vague predicate (truly) relative to *each* range of application that contains the value in question. For example, when producing the classifications in row β, you would classify patches #1 through #5 relative to each of the ranges *i-iii* (i.e., relative to *i* and relative to *ii* and relative to *iii*), patches #6 and #7 relative to each of *iv-vii*, and patches #8 through #12 relative to each of ranges *viii-x*.

What I have just proposed may sound like saying that the competent speaker has somehow to find out, independently of the classifications she makes or is disposed to make using a given vague word,

α	B	B	B	B	B	B	B	b	G	G	G	G	
β	B	B	B	B	B	b	b	G	G	G	G	G	

	1	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	
i	B	B	B	B	B								***i***
ii	B	B	B	B	B	B							***ii***
iii	B	B	B	B	B	B	B						***iii***
iv				b	b	b	b						***iv***
v					b	b	b						***v***
vi						b	b	b	b				***vi***
vii					b	b	b	b					***vii***
viii						G	G	G	G	G	G	G	***viii***
ix					G	G	G	G	G	G	G	G	***ix***
x							G	G	G	G	G	G	***x***

i	B	B	B	B	B								***i***
ii	B	B	B	B	B	B							***ii***
iii	B	B	B	B	B	B	B						***iii***
iv				b	b	b	b						***iv***
v					b	b	b						***v***
vi						b	b	b	b				***vi***
vii					b	b	b	b					***vii***
viii						G	G	G	G	G	G	G	***viii***
ix					G	G	G	G	G	G	G	G	***ix***
x							G	G	G	G	G	G	***x***

Figure 5.13. Game board. Top panel shows the rows that could be used to produce the classifications in row β. Bottom panel shows a sequence of rows that could be used in a reversal condition with hysteresis after the player has shifted to green at patch #8. On this game board the hysteresis could be up to three patches long. See www.oup.com/us/raffman for color illustrations.

whether any ranges of that word include the relevant value. But that is not what I mean. Simply in virtue of her competence with the word, her applications of it are "automatically" relativized to the appropriate ranges of application. Suppose you are looking at patch #6 on the game board; on this occasion it strikes you as blue, and so you classify it as blue. Because you are competent with 'blue,' your very act of classifying #6 automatically relativizes your application of 'blue' to range *ii* and to range *iii*. Or perhaps it would be better to say that you relativize your application of 'blue' to the relevant ranges—you use those ranges—*in the very act* of classifying the patch. You don't need to do anything else in order to relativize your judgment to the proper ranges of application. Similarly, you relativize your application of 'rich' to certain of its ranges in the very act of classifying $165,000 as rich relative to Vinny.

Let me repeat: I am not saying that vague predicates are response-dependent.[16] The responses or judgments of competent speakers do not determine whether an item satisfies a given predicate. Whether an item satisfies 'blue' relative to a given V-index is determined by the semantics of the predicate, in particular by its ranges of application relative to that V-index, and by the item's hue. An item may satisfy 'blue' relative to every range of the predicate for a given V-index, or relative to none, or relative to some but not all ranges—independently of anything to do with judging subjects or speakers. Variable (transitional) items satisfy more than one predicate: For instance, patch #6 on the game board is blue relative to range *ii* and to range *iii*, borderline relative to each of ranges *iv* through *vii*, and green relative to each of ranges *viii* and *ix*. All that competent speakers (or maybe some of their subpersonal parts) are allowed to determine is how they will classify an item, and to that extent which ranges they will use, in a particular utterance.

One might have wondered whether relativity to ranges of application is itself just a fine-grained species of V-index-relativity

(f. context-relativity). We can now see why the answer is 'no.' A V-index is a set of factors that speakers can select, that is, intentionally choose, to take into account in their application of a term. For instance, I can choose to relativize my use of 'tall' to jockeys instead of basketball players, and I can choose to take into account the number of hairs on a head rather than their arrangement, in a given application of 'bald.' Speakers are typically aware of, and can say, which V-index they are relativizing to in a given utterance. (They also typically know which V-indices their interlocutors are relativizing to and so are aware of changes of V-index during a conversation.) More important, a V-index can be, and usually is, chosen *for a reason*—for example, because it is topically relevant to a current conversation, or simply because it is the one currently in use. In contrast, as I have just explained, speakers do not *choose* or *decide* the ranges to which they will relativize their applications of vague words, and do not take those ranges *into account*. Rather, they simply decide how they will classify a given item, and relativization to (each of) the corresponding ranges of application comes along automatically, for free, as it were. Furthermore, in the variable transitional region between categories, speakers' classifications, a fortiori the ranges to which those classifications are relativized, are determined by brute psychological mechanisms. Thus they are not made *for a reason*; as I have explained, our classifications of variable items are made without (nontrivial) justification. For these reasons, among others, range-relativity is not plausibly viewed as a species of index-relativity.[17]

5.5. AN ETYMOLOGICAL SPECULATION

After working out this hysteretic solution to the dynamic sorites, I discovered some evidence of a possible etymological link between

the notions of vagueness and hysteresis. The English adjective 'vague' and cognates 'vagrant,' 'vagabond,' 'vagary,' and so on originate in the Latin 'vagus,' meaning 'wandering.'[18] (The French noun 'vague,' whose primary meaning is 'wave,' as in an ocean wave or a sound wave or, less literally, a wave of panic, has the same Latin root. See again the definition of the French adjective 'vague' that opens chapter 4.) The word 'hysteresis' is rooted through the ancient Greek *hysteros*, meaning 'late' or 'lagging behind' or 'coming after,' to the Sanskrit *ûttaras*, meaning 'uterus.' The Greek term for 'uterus,' the feminine form *hystera*, also derives from the Sanskrit word. Various theories have been offered as to the connection between these seemingly disparate strands of meaning ('late' or 'behind,' and 'uterus'). Thomas Morgan writes:

> Despite the common etymological base (in the Sanskrit *ûttaras*) of this *hysteros* and *hystera* (womb), the *Greek-English Lexicon* does not say why the womb is so called. . . . Froma Zeitlin . . . suggests that *hystera* may draw on *hysteros* to name the second, or later, sex; John Belton[,] . . . noting that one meaning of *hysteros* is 'afterbirth,' proposes that this later stage of delivery may have been thought to have been the womb itself (*hystera*), a new one of which was supposed to grow back. (1994, 54)

Mark Adair explains that the word 'uterus'

> shares a common Sanskrit origin (*uttaras*) with the adjective *hysteros*, of which one meaning is 'placed behind.' The *hystera* is behind, that is, caudal (toward the tail) to all of the other massy internal organs. (1995, 158)

Especially relevant to our present interests is the fact that the ancients conceived of the uterus as an organ, indeed a creature, that *wandered*

throughout the body, thereby causing the symptoms of hysteria. In the *Timaeus* Plato writes,

> In women…[the] womb, a living creature within them with a desire for child-bearing, if it be left long unfruitful beyond the due season, is vexed and aggrieved, and wandering throughout the body and blocking the channels of the breath, by forbidding respiration brings the sufferer to extreme distress and causes all manner of disorders.

Indeed the notion that female hysteria was caused by a 'wandering womb' persisted in European academic medicine into the Middle Ages.

Unsurprisingly, I have not seen mention of an etymological tie between 'vague' and 'hysteretic' in the scholarly literature, so I rest nothing on it. But it would be nice to think that they are distant relatives. Hysteresis is nothing if not the wandering of a system's shifting point between two discrete states.

5.6. THE TRUTH ABOUT TOLERANCE

If the major premises of both versions of the sorites paradox are false, as we believe, why do they seem true? In chapter 4 I claimed that the multi-range semantic theory points toward one possible reason why the major premise of the semantic paradox seems true. We confuse (b), which is true, with (a), which is false:

(a) The increments between adjacent items in a sorites series for vague predicate 'Φ' are sufficiently small as to leave the application of 'Φ' unaffected. In other words, the

increments are sufficiently small as to make any differential application of 'Φ' as between them *impermissible*.

(b) The increments between adjacent items in a sorites series for vague predicate 'Φ' are sufficiently small as to make any differential application of 'Φ' as between them *arbitrary*.

In fact, I think the latter confusion also accounts in part for the plausibility of the major premise of the dynamic paradox. But a second reason—the principal one, I think—for the seeming truth of the major premises of both versions of the puzzle comes to light in Part III of the experiment.

In Part III we presented subjects with consecutive pairs of neighboring patches: #1/#2, #2/#3, #3/#4, and so on.[19] The task was to classify each patch in a pair singly, as cued by a black dot. (See Figure 5.7 on p. 148.) The task was performed in the same five conditions as before (though in Part III subjects ran each condition only once). We predicted that subjects would virtually always put both members of a pair in the same category. Presumably, in order to do that, they would have to shift categories between consecutive pairs, and the shared patch in those pairs would have to be categorized differently in its two pairings. Our results confirmed this prediction, as shown in Figure 5.14. For example, in the B[r] condition the subject shifted categories from *blue* to *borderline* between the pairs #22/#21 and #21/#20. The shared patch, #21, was classified as blue in the former pair and as borderline in the latter. (Notice that hysteresis occurred also in this pairwise presentation, in the G[r] condition.)

How does this finding help to explain the intuitive appeal of the major premises of the sorites paradoxes? Consider that those premises only ever make reference to two neighboring items in a sorites series. Therefore it seems plausible that when we go to evaluate those

stm	BG			GB			Rn			B[r]					G[r]				stm
1	-	-		-	-		-	-		-	-	-	-		-	-	-	-	1
2	-	-		-	-		-	-		-	-	-	-		-	-	-	-	2
3	-	-		-	-		-	-		-	-	-	-		-	-	-	-	3
4	-	-		G	-		G	-		-	-	-	-		-	-	-	-	4
5	-	G		G	G		G	G		-	-	-	-		G	-	-	-	5
6	G	G		G	G		G	G		-	-	-	-		G	G	-	-	6
7	G	G		G	G		G	G		-	-	-	-		G	G	-	-	7
8	G	G		G	G		G	G		-	-	-	-		G	G	-	-	8
9	G	G		G	G		G	G		-	-	-	-		G	G	-	-	9
10	G	G		G	G		G	G		-	-	-	-		G	G	G	-	10
11	G	G		G	G		G	G		-	-	-	-		G	G	b	G	11
12	G	G		G	G		G	G		-	-	-	-		G	G	b	b	12
13	G	G		G	G		G	G		-	-	-	-		G	G	b	b	13
14	G	G		b	G		b	G		-	-	-	-		G	G	b	b	14
15	G	G		b	b		b	b		-	-	-	-		b	G	-	b	15
16	G	G		b	b		b	b		-	-	-	-		-	b	-	-	16
17	G	G		b	b		b	b		-	-	-	-		-	-	-	-	17
18	G	b		b	b		b	b		-	-	-	-		-	-	-	-	18
19	b	b		B	b		B	b		-	-	-	-		-	-	-	-	19
20	b	b		B	B		B	B		-	b	-	-		-	-	-	-	20
21	b	b		B	B		B	B		b	B	-	B		-	-	-	-	21
22	b	B		B	B		B	B		B	B	B	-		-	-	-	-	22
23	B	B		B	B		B	B		B	B	-	-		-	-	-	-	23
24	B	B		B	B		B	B		B	B	-	-		-	-	-	-	24
25	B	B		B	B		B	B		B	B	-	-		-	-	-	-	25
26	B	B		B	B		B	B		B	B	-	-		-	-	-	-	26
27	B	B		B	B		B	B		B	B	-	-		-	-	-	-	27
28	B	B		B	B		B	B		B	B	-	-		-	-	-	-	28
29	B	B		B	B		B	B		B	B	-	-		-	-	-	-	29
30	B	B		B	B		B	B		B	B	-	-		-	-	-	-	30
31	B	B		B	B		B	B		B	B	-	-		-	-	-	-	31
32	B	B		B	B		B	B		B	B	-	-		-	-	-	-	32
33	B	B		B	B		B	B		B	B	-	-		-	-	-	-	33
34	B	B		B	B		-	B		B	-	-	-		-	-	-	-	34
35	B	-		-	B		-	-		-	-	-	-		-	-	-	-	35
36	-	-		-	-		-	-		-	-	-	-		-	-	-	-	36
37	-	-		-	-		-	-		-	-	-	-		-	-	-	-	37

BG = stimuli presented in order (straight through) from blue to green
GB = stimuli presented in order (straight through) from green to blue
Rn = stimuli presented in random order
B[r] = reversal condition starting at blue
G[r] = reversal condition starting at green

Figure 5.14. Data from one subject for Part III of the experiment (pairwise condition). The two members of each pair were always categorized identically. For example, in the G[r] reversal condition, the subject started with the pair #5/#6 and categorized both patches as green. He then continued with green until the pair #14/#15, shifting to borderline at #15/#16. Hence the shared patch #15 was classified as green in the #14/#15 pair and as borderline in the #15/#16 pair. In addition, after reversing direction, the subject exhibited hysteresis, continuing to categorize the pairs as borderline until arriving back at the pair #10/#11, where he shifted back to green. See www.oup.com/us/raffman for color illustrations.

premises, when we reflect on their truth value, we think about only two neighboring items *pairwise*, together at the same time as it were, and mentally compare them. The results of Part III of the experiment suggest that when we do this, we always find ourselves inclined to place both items in the same category. This, I think, is the truth about tolerance: The tolerance of a vague word is a feature of its competent use with respect to incrementally different items considered pairwise—not a feature of its semantics strictly speaking. Perhaps we can say that tolerance is a pragmatic feature of vague words, embodying the commonsense precept that, all else being equal, like things should be treated alike. Maybe there is even a pragmatic rule something like this:

> (Tol) For any vague expression 'Φ,' speaker S, and items x and y considered pairwise: If x and y differ incrementally on a dimension decisive of the application of 'Φ,' then if S classifies x as Φ, S should also classify y as Φ.

No paradox is spawned by Tol, since the minor premise and conclusion of the dynamic sorites refer to items considered individually. (And, of course, the minor premise and conclusion of the semantic sorites make no reference at all to how items are or should be classified by competent speakers.)

5.7. LOOKING BACK: RULES, REASONS, AND THE GOVERNING VIEW

Many philosophers think that reference to the causal psychological mechanisms involved in language production and comprehension cannot enter into a proper understanding or theoretical account of

competent linguistic practice. For example, Louise Antony explains that, in Dummett's eyes,

> causal accounts of the mechanisms by which a speaker's language operates are simply irrelevant to an understanding of the speaker's implicit grasp of the nature of linguistic practice. Any move toward examination of the causal bases of linguistic behavior *necessarily* involves a shift from intentional to non-intentional terms of description—causal accounts of language use are thus, for him, always incommensurable with rationalistic accounts, and cannot be viewed as elaborations or supplementations of them. (1997, 198)

It is one thing to discount the significance of causal mechanisms for the semantics of vague words and quite another to discount their significance for a theory of competent use (the "practice"). According to the multiple range theory of vagueness, reference to underlying causal mechanisms provides more even than "elaboration" or "supplementation" of an account of use: As far as vague words are concerned, an account of competent linguistic practice just is in part a causal-mechanical story. It must be a partly causal-mechanical story because brute causal mechanism is what makes possible the arbitrariness essential to competent use. Brute causal mechanism is what enables competent speakers to classify incrementally different items differentially. It's not that incrementally different items cannot be classified differently, but rather that any such differential classification must be made *arbitrarily*, by a kind of psycholinguistic coin toss, if you will. And we can meet the latter requirement only if no rule specifies a particular stopping place in a sorites series. Stopping places must be determined mechanically. I said earlier that the idea of a rule dictating that we stop at no particular place may seem oxymoronic, but it is of the essence of vagueness.

Echoing Dummett, Sainsbury writes that

> any attempt to describe boundarylessness in…psychological or neurophysiological terms will…miss the normative features.…No such facts will begin to capture such aspects of the use of the word 'red' as the mandatoriness of its application in some central cases, the freedom available for borderlines, and such rules as that anything at least as red as a red thing is not merely likely to be called 'red' but ought to be so called. (1990, 263)

Sainsbury may be right about the mandatoriness of application in central cases, and the rule he mentions is correct if restricted to pairwise judgments; but he is wrong about the freedom available for borderlines (variable items, on our view).[20] In fact, I suspect that the *only* way to capture that freedom is to understand our classifications of variable items as the result of brute mechanical operations; here, freedom is arbitrariness. In addition, the smoothing effect of the psychological mechanism of hysteresis in our applications of vague words explains how a competent speaker is able to shift categories without disturbing the seamlessness of a sorites series. In the case of vague words, it is only by understanding certain aspects of the underlying machinery that we begin to understand how their semantics may be implemented in their competent use.

These remarks about the role of mechanism in a theory of vagueness bring us back, at last, to the governing view of competent language use (section 1.6). Recall that this traditional view consists of two theses (Wright, 1987): (i) competent language use is entirely rule-governed, and (ii) the rules in question are discoverable by means of various processes of rational (e.g., philosophical) self-reflection, independently of empirical investigation. Wright contends that the governing view is incoherent because those reflective

processes deliver inconsistent rules—namely, a rule that says, for example, that if two items look the same in hue, then either both are blue or neither is, and a rule that says that items that look like patch #30 are green. In contrast, I suggest that, on the one hand, the first thesis of the governing view is too stringent; while on the other hand, Wright misstates the first of the two rules he cites. In fact, *most* (not all) of our competent use of vague words is governed by at least two fully consistent rules: first, a rule that any differential categorization of incrementally different items, in particular any stopping place in a sorites series, must be arbitrary, and second, Tol, a rule that incrementally different items should be categorized identically when considered pairwise. As far as I can see, these rules were discovered by the sorts of reflective procedures endorsed in the second thesis of the governing view. Tol may have required empirical testing for its confirmation, but not for its initial discovery. Moreover, given that I just stated these rules in ordinary English, we speakers can have propositional, personal-level knowledge of them, contrary to what Wright concludes.

At the same time, competent use of a vague expression cannot be wholly rule-governed; the first thesis of the governing view is incorrect. In a certain range of cases—what we have called the 'variable' cases—the rules give out; use floats free of its semantic moorings, and the application of a vague word is, from the viewpoint of semantics, anomalous. This is the point at which mechanism enters the practice. Thus competent use is not fully determined by the semantic theory. It is not the job of a semantic theory to fully describe, a fortiori to fully prescribe, the competent use of a vague word—not even relative to a given context or V-index.

This is not to say that, with respect to the variable cases at issue, our applications of vague words are irrational or unintelligible. It's not as if we are breaking rules, or attempting to follow inconsistent rules, as some theorists have suggested (e.g., Horgan 1995b). Recognizing

that brute psychology must finally take over, Wright explains that the "(crazy) idea that competence [with a vague predicate] somehow accordingly involves *disrespecting* the rules is an artifact of a misplaced adherence to . . . an incoherent over-rationalisation of our practices" (2007, 20). (By 'over-rationalisation' Wright presumably means something like 'over-estimation of the scope of the rules of application for a vague word.'[21]) Moreover, the "purely mechanical" hysteresis observed in our experiment is a pattern or regularity in our unruly use of vague words that makes it more, rather than less, intelligible—even rational. Use wanders, but not incoherently. As Wright observes, recalling Wittgenstein, "not everything judged rationally is judged for reasons" (2007, 20).[22] Brute mechanism is what enables us to shift from one predicate to another without installing a sharp boundary in a sorites series and what enables our stopping places to be nonlegislative. Generally speaking, brute mechanism is what enables our discrete language to make contact with a seamless world.