Seconde naïveté[*]

September 29, 2016

1 Introduction and Overview

There is a very general schematic principle that a 1ary predicate 'F' might satisfy, a principle that would arguably reveal to us something very interesting about being F. If R^F is a particular (reflexive, symmetric but possibly non-transitive) relation of *closeness* along a *dimension of comparison* relevant for being F, the *no-sharp-boundaries* principle can be expressed as:

(N) For all x, y such that x R^Fs y and y is otherwise qualitatively identical with x (at least as far as dimensions of comparison relevant for being F are concerned), it is

[*]In the hermeneutic tradition, *seconde naïveté* expresses the idea of reflectively reappropriating what was once first naively lived and then critically abandoned. Epistemicism—essentially, the (genuine) negation of the major premise of the Sorites paradox—elicits incredulous stares, but, puzzlingly enough, almost no theorist who so stares dares to accept even so much as that epistemicism is false—that is, dares to accept the major premise of the Sorites paradox. In this paper, I propose to do just that by going beyond naive credulity and reflecting instead on several points of theoretical and practical progress afforded by accepting the major premise of the Sorites paradox. Earlier versions of the material in this paper have been presented in 2006 at the Arché Vagueness Seminar (University of St Andrews) and at the 3rd Navarre Vagueness Workshop in Granada; in 2015, at the LanCog Metaphysics, Epistemology, Logic and Language Seminar (University of Lisbon). I would like to thank all these audiences for very stimulating comments and discussions. Special thanks go to Zachary Barnett, María Cerezo, Richard Dietz, Patrick Greenough, Rosanna Keefe, Dan López de Sa, Sebastiano Moruzzi, Diana Raffman, Fiora Salis, Ricardo Santos, Stephen Schiffer, Célia Teixeira, Crispin Wright, David Yates and an anonymous referee. I am also grateful to the editors Otávio Bueno and Ali Abasnezhad for inviting me to contribute to this volume and for their patience throughout the process (they have experienced on their skin the claims of subsection 2.2). At different stages during the writing of the paper, I have benefitted from an AHRC Doctoral Research Fellowship, from a Jacobsen Fellowship and from the FCT Research Fellowship IF/01202/2013 on *Tolerance and Instability: The Substructure of Cognitions, Transitions and Collections* (TI), as well as from partial funds from the project CONSOLIDER-INGENIO 2010 CSD2009-00056 of the Spanish Ministry of Science and Innovation on *Philosophy of Perspectival Thoughts and Facts* (PERSP), from the project FFI2011-25626 of the Spanish Ministry of Science and Innovation on *Reference, Self-Reference and Empirical Data* and from the project FFI2012-35026 of the Spanish Ministry of Economy and Competition on *The Makings of Truth: Nature, Extent, and Applications of Truthmaking*.

not the case that [x is F and y is not F].[1,2]

For example, a plausible instances of (N) is:

(N^{high}) For all mountains x, y such that x is at most one nanometre either higher or lower than y and y is otherwise qualitatively identical with x (at least as far as dimensions of comparison relevant for being high are concerned), it is not the case that x is high and y is not high.[3]

Before proceeding further, two important remarks concerning (N) and my understanding of it should be made. Firstly, (N) has almost universally been thought to be inconsistent with the (virtually undeniable) assumption that there are both positive and negative cases of being F that, while differing along the R^F-relevant dimension of comparison, are otherwise qualitatively identical (at least as far as dimensions of comparison relevant for being F are concerned).[4] The inconsistency is supposed to be revealed by the kind

[1]The qualification 'y is otherwise qualitatively identical with x (at least as far as dimensions of comparison relevant for being F are concerned)' is intended to take care of the so frequently occurring phenomenon of *multi-dimensionality*—that is, of the fact that the correctness of the application of many predicates to an object depends on the object's location along at least two distinct dimensions of comparison. For example, the correctness of the application of 'bald' to a man depends not only on the number of hairs on his scalp, but also on their distribution, density, thickness etc. Because of multi-dimensionality, the unqualified no-sharp-boundaries principle to the effect that it is not the case that [a man is bald and a man with one more hair on his scalp is not] is, strictly speaking, false: a man with an in-principle sufficient number i of hairs, widely distributed, homogeneously dense and appropriately thick counts as not bald, whereas a man with an in-principle sufficient number $i - 1$ hairs so poorly distributed as to cover only half of his scalp, so heterogeneously dense as to leave a hairless circle in the middle and so thin as to be invisible counts as bald. Therefore, unqualified no-sharp-boundaries principles like the one just mentioned apply straightforwardly only to *uni-dimensional* predicates—in the case of *multi-dimensional* ones, a no-sharp-boundaries principle applies on a particular dimension of comparison *only under the supposition that the values on the other dimensions of comparison are held constant.* Hence the need for the qualification 'y is otherwise qualitatively identical with x (at least as far as dimensions of comparison relevant for being F are concerned)' (where 'otherwise' is also supposed to take into account possible adjustments on other dimensions of comparison required by the difference on the R^F-relevant dimension of comparison). Having noted all this, henceforth, I will usually set these "niceties" aside (scare quotes needed in view of what emerges in fn 21).

[2]At some places, I will more conveniently take (N) to have as its embedded consequence the *material implication* 'If x is F, y is F' rather than the *negated conjunction* 'It is not the case that [x is F and y is not F]' (I assume that the two are fully intersubstitutable). I will also assume that, in order to establish that material implication, it is sufficient to establish, under the supposition that x is F, that y is F (conversely, throughout, I further assume that material implication satisfies *modus ponens*). Thanks to an anonymous referee for discussion of some of these issues.

[3]Throughout, I follow the policy of referring to an instance for a specific predicate of a schematic principle by superscripting the predicate to the label of the principle (in the case of (N), sometimes leaving it to context to make clear what, in the relevant instance, R^F is supposed to be).

[4]There are, as always (hem...), exceptions (see e.g. Unger [1979]). I only note that none of the arguments I will give in favour of (N) is in the least congenial to the denial of the assumption. (Nor are they particularly congenial to denials of (N) that, instead of the real thing, offer some pale imitation, like the claim that the relevant instances of (N) are, while false, analytic (see e.g. Sorensen [2001], pp. 57–67) or the claim that they are, while false, satisfied by every contextually salient object (see e.g. Fara [2000]).) Thanks to Ricardo Santos for raising this issue.

of argument long known as '*Sorites paradox*'. I have shown in Zardini [2008a]; Zardini [2008b]; Zardini [2009]; Zardini [2014a]; Zardini [2016b] that *a satisfactory weakening of the logic exists in which* (N) *is consistent with the existence of the relevant positive and negative cases* (in a nutshell, such weakening basically consists in imposing certain kinds of restrictions on the *transitivity* of the consequence relation; see Zardini [2015a] for a general philosophical discussion of non-transitive logics). In the following, I will argue that we do use some of our predicates so as to satisfy (N)—or, at any rate, that such use is a necessary condition for achieving some important theoretical and practical purposes.[5] As a consequence of the logical point just mentioned, I will neither accept nor in fact even consider the in-principle objection to my arguments to the effect that they must be wrong [since our use of the predicates in question is consistent—or since no important theoretical or practical purpose requires an inconsistent theory—whilst (N), together with the assumption mentioned above, would breed inconsistency]. It just needn't be so.

Secondly, what I propose to do in this paper is to investigate whether giving up (N) is a viable option in our theoretical and practical life. In this regard, I should note that I'm actually much less interested in the *categorical*, *empirical* and *sociological* claim that we in effect use some of our predicates so as to satisfy (N) than I am in the *hypothetical*, *philosophical* and *normative* claim that such use would achieve important theoretical and practical purposes—or, at least, is a necessary condition for achieving such purposes. I do find the arguments to follow compelling also with respect to our actual use and will therefore argue for the stronger conclusion, but I will ultimately rest content with simply showing that such use is fully intelligible, highly valuable and hardly dispensable. Once its fine architecture and details have been brought out, I hope it will be clear that, even if it is not our actual use, it is something we should strive to incorporate into our conceptual repertoire.[6]

Clearly, (N) bears a very important relation to *vagueness*. Indeed, the *naive theory of vagueness* (defended in my works referenced in the second last paragraph) holds, roughly, that *the vagueness of a predicate consists in there not being a sharp boundary between its positive and negative cases*;[7] on this theory, then, satisfaction of (N) is what the vagueness of a predicate ultimately consists in—what the *nature* of vagueness is. Suitably developed and refined, this is the theory I favour. Admittedly, the vindication of such theory is no simple task, but a crucial part of the dialectic against its rivals consists in exposing the high costs of giving up the connection between a predicate's vagueness and satisfaction of (N).

From the point of view of the naive theory of vagueness, such satisfaction is what vagueness consists in, and so what (N) is for (the otherwise impossible achievement of

[5] As will be seen, in many cases where the need for (N) emerges, *theory and praxis are inextricably entangled* (for criticism of a different proposal connecting (N) with practical interests, see Sweeney and Zardini [2011]).

[6] Those who know me a bit will quickly recognise that some of the material to follow has certain autobiographical connotations, so that the official "important theoretical and practical purposes" might actually reflect some idiosyncratic tendencies of mine. If they do, then so be it.

[7] Thus, throughout, I use 'sharp boundary' and its like to mean the kind of thing that (N) asserts the negation of.

important theoretical and practical purposes to be presently described) is what vagueness is for. In denying that satisfaction of (N) is what vagueness consists in—indeed, in denying that such satisfaction is ever so much a necessary condition for vagueness—the rivals of the naive theory commit themselves to denying that vagueness is conducive to these important theoretical and practical purposes. Moreover, since it is very plausible that vagueness would be so conducive if anything were, the rivals commit themselves to denying that such purposes can possibly be achieved. Therefore, the more valuable these purposes can be shown to be, the less appealing the rivals will appear. This specific part of the dialectic is what this paper tries to accomplish, and, granting the success of the overarching project, what this part would then amount to is an exhibition of the grounds of the vagueness of a predicate—of what the *sources* of vagueness are. Let me stress that, even though I will identify several such sources, the discussion will by no means be meant to be exhaustive. Other sources of this complex phenomenon wait to be uncovered.

Before embarking in the details of this part of the dialectic, let me stress that it is on my view akin to the one involving the naive theory of *truth*. Even though, contrary to a highly suggestive line of thought put forward by some prominent theorists (see McGee [1991]; Tappenden [1993]; Soames [1999]; Field [2003]), I do not think that the problems concerning semantic concepts and those concerning vagueness have a common root (see Zardini [2011]), I do think that, just as a strong case can be made that nothing short of the naive theory of truth can do justice to our use of the concept of truth (and related semantic concepts such as reference, denotation, satisfaction etc.) as picking out a *universal property of representational correctness* (*not*, however, as fulfilling the function of a *universal device of disquotation*, see Zardini [2014b]; Zardini [2015b] for some discussion), an at least equally strong case can be made that nothing short of the naive theory of vagueness can do justice to many features of our use of vague concepts. This case will be set out in the following. Indeed, as will be seen, in the case of vague concepts, as opposed to the case of truth (and related semantic concepts), this strategy can be developed with respect to significantly different and apparently independent features of our use of such concepts, which partly explains why sometimes satisfaction of (N) will only be argued to be a necessary condition for a particular feature of use to achieve its purpose.

The rest of this paper is organised as follows. Section 2 shows what (N) allows us to think about the world in terms of classifications that are interesting (subsection 2.1), flexible (subsection 2.2) and easy (subsection 2.3). Section 3 shows what (N) allows us to experience of the world in terms of slow changes (subsection 3.1) and matching appearances (subsection 3.2).[8] Section 4 draws the conclusions that follow from the specific arguments given for what, according to the naive theory of vagueness, the sources of vagueness are.

[8]Each of the subsections of these two main sections thus presents a distinct line of argument in favour of (N). Also, the stark distinction between the two main sections is partly dictated by presentational needs: as will be seen, some of the arguments really pertain to both sides.

2 Thoughts Requiring the Absence of Sharp Boundaries

2.1 Substantial Differences

In this subsection, I want to expand on and generalise a suggestive line of thought first put forward by Wright [1975] (*cf* Wright [1976], pp. 231–232). In considering the high plausibility of (N) for age nouns like 'child', 'adolescent', 'adult' etc., Wright remarked that the classifications induced by such nouns "are of substantial social importance in terms of what we may appropriately expect from, and of, persons who exemplify them" (Wright [1975], p. 336). He then observed that, on the one hand, "[i]t would be irrational and unfair to base substantial distinctions of right and duty on marginal – or even non-existent – such differences" and that, on the other hand, "[o]nly if a *substantial* change is involved in the transition from childhood to adolescence can we appeal to this transition to explain substantial alterations in patterns of behaviour" (Wright [1975], p. 337). On these grounds, he concluded that, for such age nouns, "very small differences cannot be permitted to generate doubt about their application without correspondingly coming to be associated with a burden of moral and explanatory distinctions which they are too slight to convey" (Wright [1975], p. 337).

I think we can extract from this the following quite general pattern of requirements of respectively *important* and *unimportant differences* that we sometimes impose on being F. That is, sometimes:

(ID) We attach importance to being F rather than falling in some sense short of being such;[9]

(UD) We do not attach importance to small differences along the R^F-relevant dimension of comparison.

Before proceeding to apply this pattern beyond age nouns, it is important to clarify the function in this dialectical context of the phrase 'falling in some sense short of being F' and its like. In the diverse arena of contemporary theories of vagueness, there are many different ways in which x can *fall short of* being F without thereby being guaranteed to be *not* F. The range of ways of being alternative to being F is wide and well-known:

- Negation of the negation of the proposition that x is F, accompanying rejection[10] or even negation of the proposition that x is F;

[9]Focus on the (possibly) weaker requirement that we attach importance to being F rather than *not* being such and consequent failure to pay due heed to the stronger (ID) seem to me to flaw the considerations advanced in Sainsbury [1989], pp. 38–39; Sainsbury [1995], p. 28.

[10]Throughout, *rejection* (along with *denial*, its speech-act manifestation) is understood as an attitude on a par with acceptance, at least in the sense that it is not presupposed that it entails or is entailed by acceptance of the corresponding negation (see Parsons [1984]; Smiley [1996]; Tappenden [1999]; Rumfitt [2000]; Field [2003]; Priest [2006], pp. 103–115 for various arguments in favour of positing this distinctive

- Negation of the proposition that x is un$F$ (where 'un$F$' is a proximate contrary of 'F'), accompanying negation of the proposition that x is F;
- Negation of the proposition that x is either F or not F;
- Rejection of the negation of the proposition that x is F, accompanying rejection that x is F;
- Rejection that x is either F or not F;
- Negation of the proposition that it is true that x is not F, accompanying negation of the proposition that it is true that x is F (or negation of the proposition that 'x is not F' is true, accompanying negation of the proposition that 'x is F' is true), and higher-order variations thereof;
- Negation of the proposition that it is definitely (or determinately, or clearly etc.) the case that x is not F, accompanying negation of the proposition that it is definitely (or determinately, or clearly etc.) the case that x is F, and higher-order variations thereof;
- Acceptance that x is F only to some intermediary degree

and many others.

Relatedly, even better established are attempts at formulating principles that, while weaker than those licenced by (N), still try to preserve some of the intuitive force behind them, mostly by also allowing some strengthening of the original supposition that x is F. Assuming that x R^Fs y and y is otherwise qualitatively identical with x (at least as far as dimensions of comparison relevant for being F are concerned), some such principles are:

- If x is definitely (or determinately, or clearly etc.) F, y is not definitely (or not determinately, or not clearly etc.) not F (and their truth-theoretic analogues);
- If x is definitely (or determinately, or clearly etc.) F, y is F (and their truth-theoretic analogues);
- It is not both acceptable that x is F and rejectable that y is F;
- If x is F to a certain degree, y is F at least to a not significantly smaller degree.

attitude). I am actually not at all sympathetic with the philosophical reasons that typically underlie the failure of that presupposition (as they effectively involve an unlikely attempt at *psychologising* the *absolutely objective* notion of x's falling in some sense short of being F), but, for the dialectical purposes of this paper, it will do no harm and will indeed be useful to work with the notion of rejection so understood. Thanks to Ricardo Santos for pushing me on this issue.

Even setting aside the question of their (doubtful) dialectical effectiveness in preserving the spirit of (N) without its alleged paradoxical consequences (see Zardini [2013] for an argument that such principles are not so effective), such principles will be irrelevant here, our assumption being just the plain one that x is F, and our question being what follows from that with respect to y's being F.

To conclude this clarification, I will henceforth use the catch-all phrase 'falling in some sense short of being F' and its like for *every* way of being *alternative* to being F. This *alterity* between falling in some sense short of being F and being F can be understood to be such that, *metaphysically* (and so most objectively), x falls in some sense short of being F iff x *fails to be* F; *logically* (and so less objectively and more subjectively), x's falling in some sense short of being F *is inconsistent with* x *being* F; *attitudinally* (and so most subjectively), x falls in some sense short of being F iff *one should reject that* x *is* F.[11] We are now ready to expand on (ID) and (UD) following through their consequences for the truth of the relevant instances of (N) and also ready to generalise their range of application to cases beyond age nouns. We will accomplish both tasks in one go by considering cases of concepts belonging to the realms of thought (both scientific and ordinary) and action.

Thought. Suppose that x is a dog. Suppose that y differs from x at most for the fact that one of y's atoms is in a location that is within a nanometre distance from the location correlating to the location of the corresponding atom part of x. Then it should be correct to say that y is a dog as well—*anything* falling in some sense short of this would seem to draw an *arbitrary* difference between x's and y's animal status based on a difference, such as the nanometrical displacement of a single atom, that we perceive to be *irrelevant* for dogs. Let us elaborate on this. Purporting to be a biological kind, doghood is supposed to be an objectively distinguished feature of nature (pictorially, to "carve nature at its joints"). We aim at describing important biological facts by referring to instantiation of such kind by a living being rather than instantiation of anything falling in some sense short of it.[12] But there is no relevant biological joint to be carved in a nanometrical difference of one atom's location: life just doesn't go that deep. The difference between being a dog and falling in some sense short of being such is therefore *essentially coarse-grained*: *at some deeper level of reality, it requires non-minimal differences at that level*. In its thus being a *substantial* (indeed, *big*) difference, the difference between being a dog and falling in some sense short of being such is *a difference of many nanometres and many atoms rather than a difference of one single nanometre and one single atom*;[13]

[11] Notice that, if e.g. an epistemic spin in terms of clarity is given to the notion of definiteness, the relevant way alternative to being F can, for example, be understood as negation of the proposition that it is clearly the case that x is not F, accompanying negation of the proposition that it is clearly the case that x is F *and* acceptance that x is, after all, not F.

[12] Throughout, I often make use of kind- and property-talk. This is only for ease of exposition and should be considered as ultimately dispensable.

[13] Thus, every case of falling in some sense short of being a dog falls indeed "quite short" of being such. Notice that this is not understood to require the existence of a *gap* between the objects that are dogs and those that fall in some sense short being such (the whole point of shifting from the property of not being a dog to the property of falling in some short of being such was precisely to guarantee *exhaustivity*!). But it

hence, we have the *entailment from an object's being a dog and another object's falling in some sense short of being such to the two objects' not differing only by a nanometrical difference of one atom's location.* And this entailment in turn requires that, under the supposition that x is a dog, we must[14] reject any predication about y entailing that y falls in some sense short of being a dog,[15] and so we must accept that y is a dog.[16] (To counter a likely rejoinder, notice that rejection that y falls in some sense short of being a dog, accompanying rejection that y is a dog, itself counts—in a typically reflexive way—as one of the senses in which y can fall short of being a dog: recall that falling in some sense short of being F covers *every* way of being alternative to being F.) Thus, under the supposition that x is a dog, y is a dog. (N^{dog}) follows. (Alternatively, it follows from the above entailment by an extremely plausible variation on *antilogism.*)[17]

Analogously to Wright's discussion of age nouns, the irrelevance of nanometrical differences for dogs can be made to emerge also from a different angle. We may assume that x reacted in a certain way[18] to the ingestion of a certain pill *because* x is a dog.[19] But, certainly, it is not the case that x reacted in that way because x's atoms are arranged at

---

does show that the difference between being a dog and falling in some sense short of being such, although a difference of *contradictories*, shares with differences of *contraries* the feature of requiring non-minimal differences along the relevant dimension of comparison (a feature that a difference of contradictories is usually supposed not to have). The operation of falling in some sense short (and, *a fortiori*, the operation of negation) *facit saltus.*

[14]Here as well as at the other relevant places, in the strong sense of 'must' licenced by entailments, and so in the sense in which it is not the case that, under the supposition that one has a ticket of a fair lottery with, say, 1,000,000 tickets, one must accept that one will not win the lottery.

[15]Abstractly, the reasoning in question goes from φ ('x is a dog') and ψ ('y falls in some sense short of being a dog') entailing χ ('x and y do not differ only by a nanometrical difference of one atom's location') to its being the case that one must reject ψ under the supposition of φ and $\neg\chi$.

[16]Can we not then further develop this reasoning to the effect that, further supposing that z differs from y at most for the fact that one of z's atoms is in a location that is within a nanometre distance from the location correlating to the location of the corresponding atom part of y, we must also accept that z is a dog? Because what underlies the reasoning is the entailment emphasised in the last sentence in the text, whether the reasoning can be so developed will depend on the logic governing the entailment. In particular, the reasoning cannot be so developed if the logic governing the entailment is the kind of non-transitive logic referenced in section 1.

[17]Abstractly, the variation in question is that, if φ ('x is a dog') and 'Things fall in some sense short of being such that ψ' ('y falls in some sense short of being a dog') entail χ ('x and y do not differ only by a nanometrical difference of one atom's location'), φ and $\neg\chi$ entail ψ. Of course, if one is willing to accept the original version of antilogism (which I do regard as compelling even in the presence of vagueness), that allows for a simpler version of the argument to the effect that, abstractly, since φ and $\neg\psi$ entail χ, φ and $\neg\chi$ entail ψ. But, as we have seen, in this dialectical context, some theorists admit the possibility of x's falling in some sense short of being F without thereby not being F, given which the original version of antilogism becomes problematic. As far as vagueness is concerned, I myself am against such possibility, but, for the dialectical purposes of this paper, it was important to present the argument in a way that finesses that issue.

[18]Where the way in question is supposed to be one of those specified by standard biology.

[19]Throughout, the relevant explanatory claims are understood to be genuine claims of *ontological dependence*, where (ontological) dependence is a genus that comprises, possibly among other relations, *causation* and *grounding* as species. As usual, and as befits the current discussion about whether a certain kind of difference is *relevant*, I assume that such relations are *non-monotonic.* Thanks to David Yates for urging these and other clarifications on this material.

most at such-and-such distances from one another and it is not the case that the distance of even just two of them is a nanometre greater: in such context, that x falls on one side rather than the other of such an exquisite difference has *no scientific relevance*, as opposed to its falling on one side rather than the other of the difference between being a dog and falling in some sense short of being such.[20] Thus, the property of being a dog cannot be *strongly replaceable* (that is, replaceable in the *explanans* of a standard dependence claim) by any property with sharp boundaries. However, given the relations between higher-level properties and lower-level properties, the property of being a dog is strongly replaceable by a lower-level property with sharp boundaries if (and only if) it is necessarily co-extensional with it (for, if it is, the fact that something is a dog very plausibly reduces to or depends on the fact that it exemplifies the lower-level property, and so everything depending on the "former" fact can be seen as more fundamentally depending on the "latter" fact). Since what (N^{dog}) is a negation of thus implies false dependence claims, (N^{dog}) follows by *contraposition*. And, even if this extremely plausible contraposition is somehow resisted, the argument can be recast so as not to rely on it: in the context in question, there is similarly no scientific relevance in the fact that x falls on one side rather than the other of a difference that is rejected to be *any coarser-grained* than one determined by a nanometrical difference of one atom's location. The property of being a dog exhibits some sort of *emergence* with respect to properties with sharp

[20]The situation can be usefully contrasted with cases of *dependence of a higher-level property on a lower-level property*: for example, assuming that the property of being fragile depends on the property of having such-and-such structure, if it is the case that the glass broke because it was fragile, it is also the case that the glass broke because it had such-and-such structure. The situation can also be usefully contrasted with cases of *surprising cut-offs*: for example, it might be surprising but nevertheless the case that the camel's back broke because it was loaded with 37 rather than 36 bricks.

boundaries,[21] and such emergence leads to (N^{dog}).[22]

The point can be made in an equally compelling and yet tellingly slightly different way for concepts that do not purport to pick out natural kinds. Consider the concept of baldness. It too is entrenched in a sophisticated (folk) theory—about the physiological causes of baldness, the way people affected by baldness look, the social impact of baldness etc. But, on the one hand, none of these diverse aspects discriminates importantly between neighbouring numbers of hairs on one's scalp: two physiological causes such that one differs from the other only in causing the loss of just one more hair are not importantly different, one's look is not importantly altered by the addition of a single hair, one's social impact is not importantly influenced by the insertion of a simple hair. On the other hand, baldness is supposed always to affect importantly these aspects. We expect any state of falling in some sense short of being bald to be due to importantly different physiological causes than those responsible for one's being bald—otherwise, why should one change one's diet in order to avoid becoming bald? We expect a bald man to look importantly different from anyone falling in some sense short of being bald—otherwise, why should

[21] In particular, while it remains very plausible that every true principle of *supervenience* of the property of being a dog on properties with sharp boundaries is knowable *a priori* and that the effects of being a dog do *not* give rise to *downwards causation* (since essentially the same kind of judgement about false dependence claims as the one exploited in the text also indicates that "going sharp" on the effect side equally leads from truth to falsity, and so in particular that the fact that x's reaction exemplifies any lower-level property with sharp boundaries—as opposed to one of those properties specified by standard biology, see fn 18—is not caused by the fact that x is a dog), it is nevertheless the case that *there are true principles of causation concerning the former property which are irreducible to principles concerning the latter properties*. Properties without sharp boundaries live a separate causal life from that of properties with sharp boundaries. Obviously, even if the property of being a dog were not so emergent, as long as, for some reason or other, (N^{dog}) holds, that property would still be *irreducible* to and *independent* of any property with sharp boundaries (the less obvious latter because it would not be the case that something is a dog because it exemplifies a property with sharp boundaries: for one thing, by (N^{dog}), it would still be a dog even if it only exemplified a property with nanometrically different sharp boundaries). ("Wait—could a property without sharp boundaries not be reduced to or depend on a property without sharp boundaries?" "How?" "Well, taking for example the property of being bald, that could be reduced to or depend on the property of having sufficiently few hairs." ""Sufficiently" for what—maybe for replacing Yul Brynner?" "No, c'mon, ya know... for being bald!" "Ah, there you are! Doesn't look like a great reduction or dependence...") In fact, even if (N) failed to hold, the phenomenon of *multi-dimensionality* remarked on in fn 1 would arguably suffice to generate massive irreducibility and independence. For, given multi-dimensionality, for every i, the property of being bald is not necessarily co-extensional with—let alone reducible to or depending on—the property of having at most i hairs on one's scalp. Given that the dimensions of comparison relevant for the exemplification of a multi-dimensional property are typically *indefinitely many*, that presumably leaves as possible candidates for reduction or dependence only disjunctive properties with indefinitely many disjuncts each of which is a conjunction with indefinitely many conjuncts each of which states that things are such-and-such on a certain dimension of comparison. Those are horrible candidates for reduction or dependence; worse, plausibly, it is the case that a man exemplifies the candidate that is necessarily co-extensional with the property of being bald because he is bald rather than, *vice versa*, its being the case that he is bald because he exemplifies the candidate. Under the usual understanding of what kind of properties the physical level contains (which rules out higher-level properties of the kind of the property of being a dog or of the kind of the property of being bald), such emergence, irreducibility and independence would then have the *antiphysicalist* implication that *the physical level does not exhaust the fundamental level*.

[22] I owe the inspiration for this argument from dependence to a remark put to me by Stephen Schiffer.

one be relieved about one's look upon hearing a trustworthy denial that one is bald? We expect a bald man to have an importantly distinctive social impact—otherwise, why should uncertainty about one's baldness engender oscillation as to how one will be received for the first time by one's partner's parents? It follows from this contrast that baldness should not be sensitive to one-hair differences or, worse, to a nanometrical difference of one atom's location: baldness too just doesn't go that deep. The difference between being bald and falling in some sense short of being such is therefore *essentially coarse-grained: at some deeper level of reality, it requires non-minimal differences at that level.* In its thus being a *substantial* (indeed, *big*) difference, the difference between being bald and falling in some sense short of being such is a *difference of many nanometres and many atoms rather than a difference of one single nanometre and one single atom*; hence, we have the *entailment from an object's being bald and another object's falling in some sense short of being such to the two objects' not differing only by a nanometrical difference of one atom's location.* And this entailment in turn requires that, under the supposition that x is bald and y differs from x at most for the fact that one of y's atoms is in a location that is within a nanometre distance from the location correlating to the location of the corresponding atom part of x, we must reject any predication about y entailing that y falls in some sense short of being bald, and so we must accept that y is bald. Thus, under the supposition that x is bald, y is bald. (N^{bald}) follows. (Alternatively, it follows from the above entailment by the extremely plausible variation on antilogism discussed in fn 17.)

Again, the irrelevance of nanometrical differences for bald men can be made to emerge also from a different angle. We may assume that x is no longer a successful womaniser *because* x has become bald. But, certainly, it is not the case that x is no longer so because x's atoms are arranged at most at such-and-such distances from one another and it is not the case that the distance of even just two of them is a nanometre greater: in such context, that x falls on one side rather than the other of such an exquisite difference has *no ordinary relevance*, as opposed to his falling on one side rather than the other of the difference between being bald and falling in some sense short of being such. Thus, the property of being bald cannot be strongly replaceable by any property with sharp boundaries. However, given the dependence of higher-level properties on lower-level properties, the property of being bald is strongly replaceable by a lower-level property with sharp boundaries if (and only if) it is necessarily co-extensional with it. Since what (N^{bald}) is a negation of thus implies false dependence claims, (N^{bald}) follows by contraposition. And, as per the second last paragraph, the argument can be recast so as not to rely on contraposition. The property of being bald exhibits some sort of *emergence* with respect to properties with sharp boundaries, and such emergence leads to (N^{bald}).

Action. Suppose that x is a person. Suppose that y differs from x at most for the fact that one of y's atoms is in a location that is within a nanometre distance from the location correlating to the location of the corresponding atom part of x. Then it should be correct to say that y is a person as well—*anything* falling in some sense short of this would seem to draw an *invidious* difference between x's and y's personal status based on a difference, such as the nanometrical displacement of a single atom, that we perceive

to be *irrelevant* for persons. Let us elaborate on this. Personhood is at the centre of a rich web of commitments and entitlements. Importantly different actions are licenced with respect to someone who is a person and someone who falls in some sense short of being a person. And we simply haven't come up with the concept of a person to find ourselves forced to discriminate in those ways between two fellow beings differing only by a nanometrical difference of one atom's location.[23] The difference between being a person and falling in some sense short of being such is therefore *essentially coarse-grained: at some deeper level of reality, it requires non-minimal differences at that level.* In its thus being a *substantial* (indeed, *big*) difference, the difference between being a person and falling in some sense short of being such is *a difference of many nanometres and many atoms rather than a difference of one single nanometre and one single atom*; hence, we have the *entailment from an object's being a person and another object's falling in some sense short of being such to the two objects' not differing only by a nanometrical difference of one atom's location.* And this entailment in turn requires that, under the supposition that x is a person, we must reject any predication about y entailing that y falls in some sense short of being a person, and so we must accept that y is a person. Thus, under the supposition that x is a person, y is a person. (N^{person}) follows. (Alternatively, it follows from the above entailment by the extremely plausible variation on antilogism discussed in fn 17.)

Again, the irrelevance of nanometrical differences for persons can be made to emerge also from a different angle. We may assume that I should try my best to save x's life *because* x is a person. But, certainly, it is not the case that I should do so because x's atoms are arranged at most at such-and-such distances from one another and it is not the case that the distance of even just two of them is a nanometre greater: in such context, that x falls on one side rather than the other of such an exquisite difference has *no moral relevance*, as opposed to her falling on one side rather than the other of the difference between being a person and falling in some sense short of being such. Thus, the property of being a person cannot be strongly replaceable by any property with sharp boundaries. However, given the dependence of higher-level properties on lower-level properties, the property of being a person is strongly replaceable by a lower-level property with sharp

[23]Presumably, for some relevant way w, the predicate 'treated in w' (as opposed to 'to be treated in w') is precise enough so that there possibly is a finite series of the kind adumbrated in the text where an element treated in w is immediately followed by an element not treated in w (setting aside the possibilities of all the elements' being treated in w and of all the elements' not being treated in w). But that just shows that we can be forced to discriminate against our own convictions, and it is absurd to seek any safeguard against this (sadly real) possibility in features of the use of a word or a concept. It doesn't show that these convictions are wrong. We can still insist that no discrimination *should* be made on the basis of such an invidious difference, even though, when forced to deal with all the elements of the series in the same situation, we will *as a matter of fact* be forced to make such discrimination. This might well be considered a hitherto ignored aspect in which *there can be an unbridgeable gap between ought and is*, indeed a gap so wide as to generate *practical dilemmas* (since, when forced to deal with all the elements of the series in the same situation, whichever pattern we may follow in treating or not treating them in w will be wrong). The absence of sharp boundaries embodies an ideal that thus clashes with the existence of soritical totalities. (Similarly to the clash arising from the fact that, for every i, if you reasonably donate i EUR to charities, you could reasonably have donated $i + 1$ EUR instead.) Thanks to Zachary Barnett for discussion of some related issues.

boundaries if (and only if) it is necessarily co-extensional with it. Since what (N^{person}) is a negation of thus implies false dependence claims, (N^{person}) follows by contraposition. And, as per the fourth last paragraph, the argument can be recast so as not to rely on contraposition.

2.2 Stretching the Truth

In this subsection, I aim at capitalising on the phenomenon that we are willing to *stretch* the information that we gather[24] about some cases to other cases that are similar but not necessarily identical to them in the relevant respects, and that we take such stretching to be *conclusive* (in the sense that we take it that *its conclusion is guaranteed to be true if the original information is true*). If we are told that Weimar is very far from Berlin and know that Jena is very close to Weimar, we take it that we can conclusively infer that Jena is also very far from Berlin; if we are told that 10 kilometres is a very long distance to run, we take it that we can conclusively infer that 9.999 kilometres is a very long distance to run; if we are told that 10 hours of work per day is too much, we take it that we can conclusively infer that 9 hours, 59 minutes and 59 seconds of work per day is too much.[25]

Let us deepen our understanding of the phenomenon of stretching by focussing on a particular case. Suppose that I tell you that arriving at time t (specified on a second scale) is arriving roughly on time. Then it seems that you can conclusively infer that also arriving at $t + 1$ is arriving roughly on time. More generally, given that the information that arriving at t is arriving roughly on time has been gathered, it seems that one can conclusively infer that also arriving at $t + 1$ is arriving roughly on time. Interestingly, there would seem to be a restriction on the way the initial information has to be gathered. Almost any way, whether non-inferential (*via* testimony, or memory, or perception, or introspection, or intuition) or inferential, will licence the inference, *unless* it turns out ultimately to rely on an analogous inference from the information that arriving at $t - 1$ is arriving roughly on time. For suppose that I came to believe that arriving at t is arriving roughly on time just because someone else first successfully intuited and told me that arriving at $t - 1$ is arriving roughly on time and I then inferred from that that also arriving at t is arriving roughly on time. In such situation, although it seems that *my* inference is conclusive, it does not seem that *your* inference is in turn conclusive—only the unstretched truth is allowed to be stretched.[26]

[24]Throughout, I take information gathering to be a *success action*.

[25]The phenomenon is just that we take stretching to be conclusive in the sense specified in the text. Emphatically, the phenomenon is not that under no circumstances may the *prima facie* justification a subject has for the stretching be *defeated* by misleading evidence. For example, one can certainly be struck by a Sorites paradox for 'very far' in such a way as to lose one's justification for inferring 'Jena is very far from Berlin' from 'Weimar is very far from Berlin' and 'Jena is very close to Weimar', just as one can certainly be struck by the Liar paradox in such a way as to lose one's justification for inferring ' 'Snow is white' is true' from 'Snow is white'.

[26]It is natural in this respect to make a connection with certain features of the notion of *evidence* (for which see Zardini [2016a]; see also Zardini [2015a] for related discussion concerning *non-transitive logics*).

Even henceforth assuming that we are not mistaken in taking it that stretching is conclusive, we still do not have quite what we want, since ($\mathrm{N}^{roughly\ on\ time}$) does not immediately follow from the fact that, for every t, if the information that arriving at t is arriving roughly on time *has been gathered*, arriving at $t+1$ is arriving roughly on time.[27] There are at least two ways in which this gap can be bridged. Firstly, the phenomenon of stretching would seem to extend beyond information gathering so as to cover also *supposition* (substituting 'supposition' for 'original information' in the explication of 'conclusive' in the second last paragraph): given that it is supposed that arriving at t is arriving roughly on time, it seems that one can conclusively infer, under that supposition, that also arriving at $t+1$ is arriving roughly on time. Thus, under the supposition that arriving at t is arriving roughly on time, arriving at $t+1$ is arriving roughly on time. ($\mathrm{N}^{roughly\ on\ time}$) follows. Secondly, suppose that arriving at t is arriving roughly on time. Then it is certainly at least *metaphysically possible* for you to gather the information that arriving at t is arriving roughly on time (for example, someone could successfully intuit and tell you that). Because of the guarantee at issue in the conclusiveness of stretching, it follows that it is metaphysically possible that arriving at $t+1$ is arriving roughly on time, and so, since, in conceptual matters such as this, metaphysical possibility entails (relative) actuality, that it is the case that arriving at $t+1$ is arriving roughly on time. Thus, again, under the supposition that arriving at t is arriving roughly on time, arriving at $t+1$ is arriving roughly on time. Again, ($\mathrm{N}^{roughly\ on\ time}$) follows.

Two alternative, more conservative explanations of the phenomenon of stretching might seem tempting. As for the first alternative, more conservative explanation, one might think that, in the relevant cases, knowledge (and, even less plausibly, other epistemic properties such as justification) is governed by a *margin-for-error principle*, so that, for every t, knowledge that arriving at t is arriving roughly on time requires its being the case that arriving at $t+1$ is arriving roughly on time (see Williamson [1992]; Williamson [1994], pp. 216–247; Williamson [2000b], pp. 93–134 for margin-for-error principles in general and Mott [1998]; Williamson [2000a]; Sorensen [2007]; Williamson [2007]; Zardini [2012]; Zardini [2016f] for a critical discussion thereof). Given this, one could explain the conclusiveness of stretching by observing that, if the information that arriving at t is arriving roughly on time is knowledgeably gathered, it is guaranteed to be true that arriving at $t+1$ is arriving roughly on time. However, even granting for the sake of argument a margin-for-error framework, such explanation would seem problematic in several respects. To start with, it would presumably have to take the conclusive inference to be really one from 'I know that arriving at t is arriving roughly on time' to 'Arriving at $t+1$ is arriving roughly on time'.[28] But stretching also occurs in contexts where it is extremely implau-

In the following, such restriction on information gathering will be implicitly understood to be in place.

[27] Thanks to Crispin Wright for pressing this worry.

[28] An interesting variation would be to take the conclusive inference to be a transition from the *state of knowing* that arriving at t is arriving roughly on time to the *state of believing* that arriving at $t+1$ is arriving roughly on time. Contrary to a natural tendency in contemporary theorising, I am actually in general quite happy with acknowledging irreducibly *state-based* inferences (in a suitably broad sense of 'state') over and above *content-based* inferences (see Zardini [2016f]); however, I do not think that the variation in question offers sizeable advantages over the version considered in the text. For it avoids the

sible to suppose that the subject is engaged in such higher-order reasoning in the first place (for example, she might simply lack the concept of knowledge), let alone that she has the higher-order knowledge plausibly required for the inference to confer knowledge of its conclusion (for example, she might simply be a sceptic about knowledge). Moreover, stretching seems no less conclusive when the information is gathered unknowledgeably outwith the putative margin for error—what triggers conclusiveness would seem to be the plain information gathering rather than its knowledgeability.[29] Finally, in the case of stretching from supposition rather than from information gathering, the question of knowledge or of margin-for-error principles does not even arise.

As for the second alternative, more conservative explanation, one might think that the phenomenon can be explained in terms of the *high, though usually* < 1 *epistemic probability of truth preservation*[30] by these inferences (see Sorensen [2001], pp. 57–67 for an (N)-unfriendly perspective on the positive epistemic properties of these inferences). The phenomenon would then be revealed to be essentially of the same kind as the inference from someone's having a ticket of a fair lottery with, say, 1,000,000 tickets to her losing the lottery: this inference too enjoys a very high epistemic probability of truth preservation which plausibly makes it legitimate.[31] However, such model would not seem to fit all the aspects of the phenomenon of stretching. In particular, the non-conclusiveness of the (highly probably truth preserving) inference from someone's having a ticket of the lottery to her losing the lottery would seem completely lacking in the case of stretching inferences. In turn, this difference is reflected in the fact that the latter inferences *support the universal closure of their material or non-material conditionalisations* (their epistemic probability of truth preservation would not seem considerably higher than the epistemic probability of 'For every time t, if arriving at t is arriving roughly on time, arriving at $t+1$ is arriving roughly on time') whereas the former do not (their epistemic probability of truth preservation is considerably higher than the epistemic probability of 'For every subject s, if s has a ticket of the lottery, s will lose the lottery'). The difference is further reflected in the fact that the latter inferences *support their non-material conditionalisations* (their epistemic probability of truth preservation would not seem noticeably higher than the epistemic probability of any instance of the non-material conditional 'If arriving at t is arriving roughly on time, arriving at $t + 1$ is arriving roughly on time') whereas the

first problem to be raised in the text only to plunge headlong into a commitment to the rebarbative claim that, sometimes, the conclusive inference cannot confer knowledge of its conclusion (on pain of validating ($\mathrm{N}^{\textit{known to be roughly on time}}$)), and it is just as subject to the second and third problems to be raised in the text.

[29] I am very plausibly assuming that information gathering does not entail knowledge. In general, just as I do not think that every factive mental state entails knowledge (see Zardini [2016e]), I also do not think that every success action does so: for example, the action denoted by Italian *indovinare che l'Italia vincerà il prossimo Mondiale* is a success action but does not entail knowledge that Italy will win the next World Cup.

[30] Throughout, *pace* e.g. Read [2003], I naturally understand truth preservation in terms of *material* implication.

[31] A conspicuous group of philosophers (see e.g. Nelkin [2000]) would disagree about this particular kind of example. But the points I will be making apply with similar force to many other less controversial kinds of examples (for instance, the inference from the wall's looking blue to its being blue).

former do not (their epistemic probability of truth preservation would seem noticeably higher than the epistemic probability of any instance of the non-material conditional 'If s has a ticket of the lottery, s will lose the lottery'). The difference is yet further reflected in the fact that the latter inferences *are still compelling when reasoning suppositionally rather than categorically* whereas the former are not.[32]

Once the existence of stretching inferences requiring the relevant instances of (N) has been established, it is instructive to investigate which purposes such inferences serve. Though possibly very *similar*, most things are not *exactly* alike in any respect. Gathering information about an object, we should wish to be in a position to apply it to every object relevantly similar to it (even if not as good as the original object along the relevant dimensions of comparison): our extrapolations about the world would otherwise be seriously hampered, for the new objects encountered are very seldom exactly similar in some respect to the old ones (and often not as good as the old ones along the relevant dimensions of comparison). Hence, it would seem that there is a strong natural pressure towards adopting concepts satisfying (N): for there certainly is a strong natural pressure for gathering information in such a way as to make the widest possible use of it in the presence of new, almost certainly not exactly similar cases,[33] and the only way to achieve this is to conceptualise the information with concepts satisfying (N). Thus, to return to our case, there is a typical complex cluster of pieces of information that one can gather about a certain time t, which, as a first stab, could be expressed by saying that arriving at t would not spoil the point of meeting, that by arriving at t one would not fail to comply with one's commitments, that people would not be annoyed were one to arrive at t etc. We should wish to be in a position to apply this wealth of information to every time relevantly similar to t (even if slightly later than it), and this is only possible if the several pieces of information are collected together into a concept (like the concept of being a time such that arriving at it is arriving roughly on time) that satisfies (N).

Another at least equally important source of the need for stretching comes most clearly into view by reflection on the point of so-called "*vaguefiers*" in natural languages (see Lewis [1970]; Lakoff [1973]; Zadeh [1975]; Kamp [1975] for discussions and theories of vaguefication). One calculates and thinks that one can make it to the date by t, but one soon realises that likely slightly delaying circumstances may occur or that likely subtle miscalculations might have occurred. A rough-and-ready calculation shows that δ is a reasonable margin for error. The question then arises as to why one does not rest content

[32] If we wish to locate stretching inferences in the insightful taxonomy of Salmon [1967], pp. 5–11, in being *conclusive* stretching inferences are *demonstrative* and, yet, in inferring conclusions that are in an intuitive sense *stronger* than their premises, they are also *ampliative* (and are so in a sense similar to the one in which inductive inferences are ampliative, a sense that goes beyond all the senses usually envisaged for how a demonstrative inference can be ampliative; see also Zardini [2015a], pp. 263–264). This reveals one aspect in which, contrary to the refrain linking vagueness with ignorance, in being conducive to achieving certain important theoretical and practical purposes, concepts without sharp boundaries actually also increase the quantity and improve the quality of our knowledge.

[33] Of course, the *width-of-use* requirement has to strike a balance with a contrary, but equally pressing, *informativity* requirement: while we want the information to be applicable to *enough many* cases, we don't want it to be applicable to *too many* of them (or, worse still, to *all* of them!).

with the promise of arriving *between* t *and* $t + \delta$, rather than vaguefying and promising to arrive *at about* t?

One minor reason is that any particular bounded interval would seem completely arbitrary, triggering undesired conversational *implicatures* ("Why did he choose exactly 'at $t+\delta$' rather than 'at $t+\delta+0.000001$'?" would ask herself the heedful lover). However, use of predicates satisfying (N) is by itself neither necessary nor sufficient to dispose of the arbitrariness induced by predicates that do not do so. It is not necessary because arbitrariness can be lost by *underspecification* rather than vaguefication, by choosing an appropriately coarse-grained classification of times (so that one can promise to arrive at 8 pm, meaning by this to arrive at any time between the beginning and the end of the hour). It is not sufficient because arbitrariness still attaches to 'at about t' ("Why did he choose exactly 'at about t' rather than 'at about $t + 0.000001$'?" would ask herself the heedful lover) as long as 't' designates a sufficiently specific time (see Alston [1964], p. 85; Sorensen [1989]; Burns [1995] for discussions of the relation between vagueness and underspecificity).

The major reason for preferring vaguefication lies elsewhere and can be seen as follows. Two contrasting requirements on fixing a time for a date can be identified:

(a) In fixing a time, one commits oneself to be there at that time. This circumstance pushes towards a generously late time—one tends to make one's life as easy as possible;

(b) However, of course, too late a time may put into jeopardy some if not all the purposes of the date. This other circumstance pushes towards a not too generously late time—life is never too easy.

As a matter of fact about our ordinary circumstances, δ will therefore be bound not to take into account all the likely slightly delaying circumstances that may occur or all the likely subtle miscalculations that might have occurred. But *this puts any candidate for being* δ *that is good enough at steering a middle course between the two contrasting requirements (a) and (b) under a terrible pressure*: for, if as many delaying circumstances occur or subtle miscalculations have occurred as are allowed by the candidate (a possibility that is after all in the very essence of the candidate to take seriously enough!), the occurrence of *one single additional* slightly delaying circumstance or of *one single additional* subtle miscalculation will suffice to engender a failure to comply with one's commitment. Clearly, shift to the successor candidate will do little to alleviate this pressure, and even that little is likely to be offset by a lower score on the dimension of requirement (b). What is needed to remove the pressure is of course a specification of the time *which already includes a provision for slightly later times* (those that would be needed should one single additional slightly delaying circumstance occur or should one single additional subtle miscalculation have occurred)—what is needed is a specification of the time with a concept satisfying (N).[34]

[34]Notice that underspecification instead of vaguefication will not do it, for, as a matter of fact about our ordinary circumstances, the coarse-grained time chosen will equally be bound not to take into account

The kind of situation envisaged can be modelled in the following way. For simple enough cases, we can assume that the set X of all such likely-to-happen impediments is such that all of its members enjoy more or less the same (high) epistemic probability and are independent from one another. δ can then be seen as allowing for a certain finite number i of such events to happen. Since δ is a best candidate, requirement (a) implies that the epistemic probability that all the members $x_0, x_1, x_2 \ldots x_{i-1}$ of any $Y \subseteq X$ whose cardinality is i unluckily happen together is not too high, while requirement (b) jointly with the facts of the matter about our ordinary circumstances implies that it will not be as low as one might ideally wish. In many cases, this will arguably constrain the value of the epistemic probability of all the members' of any such Y unluckily happening together in the neighbourhood of .2 (at least according to my personal estimate of the pros and cons!). But then there will be a .2 probability that one only makes it *just* on time—a .2 probability that a situation will be realised where the highly epistemically probable happening of *one single additional* impediment will tilt the balance from complying with one's commitment to failing to comply with one's commitment. Even if low, a value of .2 still represents an unreasonably high risk of being in a situation where one is very likely to be subject to the cruel mockery of failing to comply with one's commitment because of the occurrence of a single, minute, in itself insignificant impediment. The shift to a concept satisfying (N) avoids this, since the time interval selected will now be such that, if it allows for a certain finite number i of impediments to happen, it also allows for $i + 1$ such events to happen. This does not of course guarantee that one will arrive on time, nor that the epistemic probability of this not happening is as low as one might ideally wish (in the model just sketched, it can e.g. be set to be only slightly lower than .2), but it does ensure that the epistemic probability that one will be in a situation where one is very likely to be subject to a cruel mockery is 0.

The major reason for preferring vaguefication in this case has thus been traced to the need of insuring oneself from failing to comply with one's commitment due to the occurrence of a single, minute, in itself insignificant impediment. Such need can be satisfied only if the commitment is expressed using a concept that allows for stretching. Having worked out the major reason for vaguefication with respect to a very particular case, it is easy to see how this reason can be generalised to a wide range of other cases. For many F, two contrasting requirements on accepting that x falls under the concept of being F can be identified:

(a′) In accepting that x falls under the concept of being F, one commits oneself to x's being good enough as to meet a sufficient condition for falling under the concept of being F. This circumstance pushes towards generously weak sufficient conditions for falling under the concept of being F—one tends to make one's life as easy as possible;

(b′) However, of course, too relaxed sufficient conditions may put into jeopardy some if not all the purposes of applying the concept of being F. This other circumstance pushes towards not too generously weak sufficient conditions—life is never too easy.

all the likely slightly delaying circumstances that may occur or all the likely subtle miscalculations that might have occurred.

As a matter of fact about our ordinary circumstances, any property with sharp boundaries candidate for being picked out by the concept of being F will therefore be bound not to take into account all of x's likely deviations along any dimension of comparison relevant for being F. But *this puts any such candidate that is good enough at steering a middle course between the two contrasting requirements (a′) and (b′) under a terrible pressure*: for, if as many deviations occur as are allowed by the candidate (a possibility that is after all in the very essence of the candidate to take seriously enough!), the occurrence of *one single additional* deviation will suffice to engender a mistake of some kind or other in acceptance that x is F. Clearly, shift to a slightly more generous candidate will do little to alleviate this pressure, and even that little is likely to be offset by a lower score on the dimension of requirement (b′). What is needed to remove the pressure is of course a specification of the sufficient conditions for falling under the concept of being F *which already includes a provision for slightly weaker sufficient conditions* (those that would still be satisfied should one single additional deviation occur)—what is needed is a concept satisfying (N).

I would like to close the discussion of vaguefication by proposing a new argument for (N) involving *hedging* vaguefiers.[35] On the one hand, it is extremely plausible that at least some hedging vaguefiers have a *finite* hedging power (indeed, a *non-iterable* hedging power). For example, '6 ft tall' can be hedged by 'roughly 6 ft tall', but the idea that the latter can in turn be hedged by 'roughly roughly 6 ft tall' is extremely dubious. We—or at least I—would not seem to have a conception of a height that, whilst good enough to count as being roughly roughly 6 ft tall, is not good enough to count as being roughly 6 ft tall. In an intuitive sense of the word, 'roughly 6 ft tall' "fuzzified" the area sharply demarcated by '6 ft tall': what else remains for 'roughly roughly 6 ft tall' to do? We—or at least I—would not seem to have a conception of a non-trivial fuzzification of a fuzzy area. (If you do, as I have already intimated, I doubt you also do for every finite iteration of 'roughly'.) Considering:

($\text{ROUGH}_{\Downarrow i}$) For every x, if x is roughly roughly roughly ... roughly (i times) F, x is roughly roughly roughly ... roughly ($i-1$ times) F,

we—or at least I—are thus led to ($\text{ROUGH}^{6\ ft\ tall}_{\Downarrow 2}$).[36] On the other hand, 'roughly' would always seem to obliterate the existence of small differences:

[35]Not all vaguefiers have hedging effects: for example, 'extremely' vaguefies precise adjectives such as 'acute' (as applied to angles), but, clearly, 'extremely acute' does not hedge 'acute'!

[36]The issues surrounding ($\text{ROUGH}_{\Downarrow i}$) have a formal parallel in the debate about *high-order vagueness*, where one interesting question is whether the analogues of ($\text{ROUGH}_{\Downarrow i}$) for 'not definitely not' hold (help: under extremely plausible assumptions, that boils down to the question whether 'definitely' iterates at $i-1$). That the question is formally parallel should however not hide important relevant differences between 'roughly' and 'not definitely not' even when restricting to (gradable) adjectives. While 'roughly' is a vaguefier that broadens the extension of precise adjectives but that would seem to leave unaffected the extension of vague ones, 'not definitely not' is on the contrary a *precisifier* that broadens the extension of vague adjectives but would seem to leave unaffected the extension of precise ones. However, in spite of its precisifying function, 'not definitely not' would not seem able to eliminate vagueness *completely*, so that, if 'F' is vague, 'not definitely not F' would also seem vague (albeit to a lesser extent). This

(ROUGH$_{\Rightarrow}$) For all x, y such that x R^Fs y, if x is F, y is roughly F.

However, given (ROUGH$^{6\ ft\ tall}_{\Downarrow 2}$) and (ROUGH$^{roughly\ 6\ ft\ tall}_{\Rightarrow}$), (N$^{roughly\ 6\ ft\ tall}$) follows.

The argument arguably extends beyond relatively *recherché* 'roughly'-predicates to cover run-of-the-mill vague predicates. For example, the idea that 'tall' can be hedged by 'roughly tall' is extremely dubious. We—or at least I—would not seem to have a conception of a height that, whilst good enough to count as being roughly tall, is not good enough to count as being tall.[37] In an intuitive sense of the word, 'tall' is associated with a "fuzzy" area: what else remains for 'roughly tall' to do? We—or at least I—would not seem to have a conception of a non-trivial fuzzification of a fuzzy area. However, given (ROUGH$^{tall}_{\Downarrow 1}$) and (ROUGH$^{tall}_{\Rightarrow}$), (N^{tall}) follows.[38]

2.3 Telling Just by Looking

In this subsection, I will work from another remark made by Wright [1975] (*cf* Wright [1976], pp. 230–231). In considering the high plausibility of (N) for a noun like 'heap', Wright briefly remarked that "'[h]eap' is essentially a coarse predicate, whose application is a matter of rough and ready judgement [...] [i]t 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 [...] our conception of the conditions which justify calling something a heap of sand is such that the justice of the description will be unaffected by any change which cannot be detected by *casual* observation" (Wright [1975], p. 335).

combination of features grounds the expectation that each new iteration of 'not definitely not' leads to a further broadening of the extension, so that the analogues of (ROUGH$_{\Downarrow i}$) for 'not definitely not' fail to hold. However, the expectation is not really conclusively confirmed by intuition, and there are in fact theoretical reasons for thinking that some such analogue does hold (see Zardini [2006a]; Zardini [2006b]; as these papers show, under extremely plausible assumptions, the analogue would then lead to (N) even in the absence of the analogue of (ROUGH$_{\Rightarrow}$) below in the text for 'not definitely not', but then again, as Zardini [2013] shows, that point does not speak against those reasons, since, under very similar assumptions, paradox ensues anyways). Be that as it may, understanding definiteness as a kind of determinacy, the analogue of (ROUGH$_{\Rightarrow}$) for 'not definitely not' is simply untenable in most logics of vagueness (given that its contrapositive, together with the compelling rule from φ to 'It is determinately the case that φ', would then licence the inference from 'y is definitely not F' to 'x is definitely not F'). It is somewhat more plausible if definiteness is understood as a kind of clarity, but then there are also somewhat stronger reasons for rejecting the analogues of (ROUGH$_{\Downarrow i}$) for 'not definitely not'. In sum, there would not seem to be similarly straightforward and compelling reasons for joint acceptance of both some of the analogues of (ROUGH$_{\Downarrow i}$) and the analogue of (ROUGH$_{\Rightarrow}$) for 'not definitely not'. Thanks to an anonymous referee for suggesting the parallel with higher-order vagueness.

[37] Since 'roughly' is typically used as a non-trivial vaguefier, one might have some initial, automatic inclination to parse 'roughly tall' as having a broader extension than unmodified 'tall' does. But I take it that such inclination quickly vanishes as soon as one reflects on the vague extension of 'tall'. Indeed, the pronounced scarcity of occurrences of 'roughly tall' is evidence that it is not felt to have a different extension from that of the more concise 'tall'.

[38] The broad kind of vaguefication performed by 'roughly' is not the only one affording the argument in the text. For example, one of my favourite variations of the argument uses 'almost' instead.

As it stands, I think that, suggestive as it may be, this remark is in need of crucial supplementation. For what it only shows is the desirability of predicates whose application can be decided by looking[39] *for some cases*—that this property does not force satisfaction of (N) can be seen by reflecting that it is also possessed by predicates like 'within a 3.152012 metre distance' which uncontroversially do not satisfy (N) (see Sainsbury [1995], pp. 27–28; Weintraub [2004], pp. 237–238). The question then naturally arises whether there is any reason relating to applicability by looking which would lead us to use a predicate that possibly satisfies (N) (like 'close') rather than one that does not (like 'within a 3.152012 metre distance').

I divide the argument in favour of a positive answer to the foregoing question in two legs. Let us call 'an *occasion*' any situation with respect to which a predicate is applied and let us call 'a *case*' any object to which a predicate is applied. Then the first leg of the argument reflects on the fact that, for many predicates, we do assume that, *on many occasions*, for *every* easily possible case,[40] under epistemically normal conditions,[41] *were* it to occur it *would* be possible to decide its application by looking. For example, we do assume that, on many occasions, for every easily possible distance, were an object located at that distance it would be possible to know by looking whether that object is close or not.[42]

Consider for instance the command:

(C_0) Slow down just in case an animal is close to the racing track!

On many occasions, (C_0) does not strike us as far-fetched at all (indeed, it is often issued!), very much in contrast with the command:

(C_1) Slow down just in case an animal is within a 3.152012 metre distance from the racing track!

Arguably, (C_1) strikes us as far-fetched because whether an object is within a 3.152012 metre distance is not the kind of question that is guaranteed to be decidable by looking when one is zooming on the racing track. If that is so, presumably (C_0) does not strike us as far-fetched because whether an object is close is the kind of question that is guaranteed to be decidable by looking when one is zooming on the racing track. But that a *guarantee* is present in the latter but not in the former case would seem to require that the *weak observationality* principle:

[39] In a suitably broad sense of 'looking', which should include use of any of the sensory modalities.

[40] An *easily possible case* is a case that could easily have occurred. *Easily*, but not necessarily *likely*, similarly to how a ticket of a fair large lottery could easily, but not likely, win.

[41] This qualification is meant to set aside any possible epistemically disturbing factors such as the lighting's being weird, the subject's being drunk, there being fakes around etc. (see e.g. Zardini [2016c] for discussion of some of such factors); in the following, I will leave it implicit.

[42] Throughout, I assume that the only dimension of comparison relevant for being close is distance.

(WOBS) On many occasions,[43] for every easily possible case x, were x to occur it would be possible to know by looking whether x is F

holds in the latter but fails to hold in the former case.[44]

Notice that (WOBS$^{\textit{within a 3.152012 metre distance}}$) does not hold because, even if, on an occasion, there are *in fact* only easy positive and negative cases of being within a 3.152012 metre distance, on no occasion, for every easily possible case x, were x to occur it would be possible to know by looking whether x is within a 3.152012 metre distance, since, for example, there are easily possible cases of objects within a 3.152012 metre distance but not within a 3.152011 metre distance, but, for no such cases, were they to occur, it would be possible to know by looking whether they are within a 3.152012 metre distance. Notice also that, since objects at *any* (reasonable) distance are on at least some of the occasions covered by (WOBS^{close}) easily possible cases (recall fn 40), and since, for every easily possible case on an occasion, there is an occasion where that case does occur, (WOBS^{close}) implies the rather strong knowability principle that, for every distance, on some occasion, it is possible to know by looking whether an object at that distance is close (henceforth, (WOBS^{close}) will be understood as carrying this implication).

Let us assume that, in some sense or other, there are *borderline cases* of 'close'. If such borderline cases constituted a range of unknowable cases, (C_0) should strike us just as far-fetched as (C_1). The surprising fact is, however, that it doesn't. Wait! Isn't it the case that we are not struck because we are in some sense *ignoring* the possibility of borderline cases? Until clear independent evidence has been presented for the postulation of a mechanism that should trigger the ignoring, such suggestion cannot be adequately discussed. But it's hard to believe that there is some such mechanism for the borderline cases of 'close' but not for the hard cases of 'within a 3.152012 metre distance', as the suggestion would require if it is to be compatible with the asymmetry of our reactions to (C_0) and (C_1). Moreover, a supplementation of (C_0) with the command of taking a picture of any animal within range δ (where δ is included in the range constituted by the borderline cases, and where the command is so understood as allowing for "playing safe" and taking pictures of animals outwith δ) does not seem vacuous. If the postulated mechanism is subtle enough as to accommodate for this apparent non-vacuity, the surprising fact that (C_0) as so supplemented still does not strike us as far-fetched can then be taken as the evidence favouring (WOBS^{close}).[45]

[43] I will chiefly conduct the discussion in terms of the stronger 'many' because it is justified by the considerations adduced in this subsection, even if all the argument really needs is the weaker 'some'.

[44] It may be tempting to try to be more conservative and weaken (WOBS) by substituting something along the lines of 'many' for 'every' (after all, a few philosophical errors have arguably been induced by obsession with the universal quantifier, see Zardini [2016d] discussion of some prominent examples). But the resulting principle presumably holds for 'within a 3.152012 metre distance', thereby failing to explain the asymmetry of our reactions to (C_0) and (C_1).

[45] Of course, this is not to deny that some cases (like the borderline cases) might be *harder than others* to decide—after all, this is so for virtually every predicate of a natural language, and 'close' is no exception. What our acquiescing reaction to (C_0) does show however is that, hard as they may be, such cases too are understood not to lie beyond what one can know just by looking. To set aside any blurriness in intuition

The second leg of the argument reflects on the fact that the *strong indiscriminability* principle:

(SIND) For all x, y such that x R^Fs y, it is not the case that, [on some occasion, it is possible to know by looking that x is F and, on some occasion, it is possible to know by looking that y is not F]

hardly needs any extended argumentation in the case of 'close'.

Putting the two legs of the argument together, we can conclude that only if (N^{close}) holds can 'close' be applied by looking in the sense required by (WOBS^{close}). For suppose that x is only a nanometre closer than y is. By (WOBS^{close}), if x is close and y is not close, [on some occasion, it is possible to know by looking that x is close and, on some occasion, it is possible to know by looking that y is not close], and so, by an extremely plausible contraposition, if the latter is not the case, neither is the case that x is close and y is not close. By (SIND^{close}), (N^{close}) follows.

It is tempting to try to reach the same result by appealing, instead of (SIND), to the *weak indiscriminability* principle:

(WIND) For all x, y such that x R^Fs y, it is not the case that, on some occasion, it is possible to know by looking that x is F and possible to know by looking that y is not F.

Emphatically, (WIND) does not entail (SIND). This is so because, although, given that x R^Fs y, it might be the case that, [on some occasion, it is possible to know by looking that x is F and, on some occasion, it is possible to know by looking that y is not F], it might still not be the case that those two epistemic feats are possible *on the same occasion*: for example, on every occasion, one's *prima facie* justification for believing that x is close may be defeated by one's *prima facie* justification for believing that y is not close.

Crucially, we cannot hope to run the previous argument by only relying on (WIND) as our indiscriminability principle while appealing instead to a stronger observationality principle, like the *strong observationality* principle:

(SOBS) On every occasion, for every easily possible case x, were x to occur it would be possible to know by looking whether x is F.[46]

---

that may be induced by the counterfactual conditional, suppose that there is in fact an animal within δ. The surprising fact is that (C_0) still does not strike us as far-fetched.

[46]The argument would then run thus. Suppose that x is only a nanometre closer than y is. Since, on some occasion, both x and y occur, by (SOBS^{close}), if x is close and y is not close, on some occasion, it is possible to know by looking that x is close and possible to know by looking that y is not close, and so, by an extremely plausible contraposition, if the latter is not the case, neither is the case that x is close and y is not close. By (WIND^{close}), (N^{close}) would follow.

We cannot do so because the extra strength of (SOBS) as opposed to (WOBS) is not supported by the previous considerations concerning the asymmetry of our reactions to (C_0) and (C_1)—the asymmetry has been observed to occur only on *many* occasions, not necessarily on *every* occasion. Moreover, the extra strength of (SOBS) is not only unmotivated, but is in itself questionable. Suppose that there is a suitable series of animals approaching the racing track each of which is suitably close to its immediate neighbours along the direction orthogonal to the racing track (and, to sidestep issues of visibility, suitably far from them along the direction parallel to the racing track). The surprising fact is that (C_0) now does strike us as just far-fetched as (C_1).[47] The impossibility of running a (SOBS)/(WIND) argument instead of a (WOBS)/(SIND) argument should not however distress us too much, given that the uncontroversial *limits* of our discriminatory powers very plausibly sustain not just (WIND), but also (SIND) (see Wright [1987], pp. 239–243; Fara [2001], pp. 916–920 for some discussion on what these limits exactly entail).

The foregoing example hints at the point of having predicates satisfying (WOBS). Most tasks involve *qualified command* for an action: the action should be performed if (qualified *command*) but also only if (*qualified* command) a certain condition obtains. For some such tasks, any condition would do, as long as its obtaining is at least in principle ascertainable (in the sense that, for every easily possible case x, a subject could sooner or later arrive at the correct answer to the question whether x satisfies the condition). However, at least for *limited* beings like us, many tasks involve *time constraints*, and for some of them these are such that the subject in charge of the task can only afford to have a *quick look* at her environment. As the asymmetry of our reactions to (C_0) and (C_1) witnesses, these latter tasks would rightly strike us as far-fetched if expressed with predicates failing to satisfy (WOBS), for they would then require a subject to *do something there is no guarantee she can do*. Predicates satisfying (N) find a source of their usefulness exactly in this kind of situation, as satisfaction of (N) by a predicate is a consequence of the predicate satisfying (WOBS) and (SIND).[48] If there has ever been a guarantee that someone can do something on the fly, the relevant instances of (WOBS) and so of (N) must hold.[49] It might be worth noting that one kind of task involving qualified

[47]The contrast with the occasion described in fn 45 is stark (and telling). Notice that, for limited beings like us, a *guarantee* can reasonably be understood as a guarantee *merely* vis-à-vis *easy possibilities*. However, since these are *not closed under conjunction*, the fact that, on many occasions, for every distance, it is an easy possibility that an object at that distance is present does not imply that, on such occasions, it is an easy possibility that, for every distance, an object at that distance is present. Therefore, on many occasions, the occasion just described in the text can be regarded as not being an easy possibility (contrary to the possibility that an object at an arbitrary distance is present), and so as not affecting the guarantee that it is knowable by looking whether an object is close. Notice also that I do not mean to imply that it is not the case that, on the occasion just described in the text, one can know by looking of each case whether it is close or not (see Zardini [2012] for a way of understanding how such knowledge is possible); the point is merely that our reaction to (C_0) on this occasion does not provide evidence for that claim.

[48]Of course, (WOBS) and (SIND) are also satisfied by precise predicates that are true of either everything or nothing in the relevant domain, but most qualified commands require distinctions that cannot be drawn by such predicates.

[49]This "*transcendental*" argument in favour of the existence of predicates satisfying (WOBS) as a

command is the one of *describing the relevant bits of one's environment employing a certain range of predicates*: this often involves *time constraints* and sometimes these are such that the subject in charge of the task can only afford to have a *quick look* at her environment. As an application of the point just made, in the latter cases the subject had better employ predicates satisfying (WOBS). If there has ever been a guarantee that someone can describe something on the fly, the relevant instances of (WOBS) and so of (N) must hold.[50]

3 Experiences Requiring the Absence of Sharp Boundaries

3.1 Seamless Change

In this subsection, I want to focus on what is probably one of the most basic ways in which the absence of sharp boundaries presents itself to a sentient creature capable of conceptualising her experience.[51] Some changes from red to orange are *seamless*. That is, in such changes, the object's change from being red to being orange seems to *take time*,

condition of possibility of on-the-fly action is confirmed by the "*empirical*" observation that there seem indeed to be features of things which are both *manifest* and *apprehensible* in certain looks. To take a prominent example, on many occasions, for every easily possible case x, were x to occur and be blue x would look blue to the relevant subject; if all this were so, the subject would see that x is blue; if all this were so, the subject could tell and so know that x is blue. (I should here note that I do not accept any straightforward implication from seeing to knowing (see Zardini [2016c]). I regard the third step in the previous train of thought as warranted rather by the fact that concepts like that of blue, by either occurring in the contents of experience or by being immediately used in conceptualising the contents of experience, *are designed to make available the contents of experience for other attitudes like belief* (the ambivalence that many natural languages register in their words relating to looks between a "phenomenal" and a "doxastic" sense is suggestive—and this very sentence hints at three examples thereof!): one believes that x is blue on the basis of one's seeing that x is blue not by *discriminating* what it is that one sees, but by simply *redeploying* at the level of belief what one sees (as suggested anyways by the idea that one "takes at face value" what one sees; the parallel with knowledge of externalistically individuated contents is telling).) The other, negative component of (WOBS^{blue}) (for the case where x is not blue) is derivable from the positive components of the instances of (WOBS) for the other colour predicates (given *exhaustivity* of colours and the relevant subject's knowledge of *exclusivity* of colours).

[50]Given the sweeping range of easily possible cases, innocent-looking qualified commands with time constraints thus turn out to carry the heavy-weight presumption (embodied in (WOBS)) that we can know by looking of every possible case whether it falls on one side or the other of a certain distinction; such presumption can actually be vindicated in spite of the limits of our discriminatory powers (codified by (SIND)), but those limits leave their mark by forcing the operative distinction to be one satisfying (N). By eliminating the boundaries on which objects are F, (N) thus removes the limits of which objects we can know by looking to be F. This reveals another aspect (additional to the one mentioned in fn 32) in which, contrary to the refrain linking vagueness with ignorance, in being conducive to achieving certain important theoretical and practical purposes, concepts without sharp boundaries actually also increase the quantity and improve the quality of our knowledge.

[51]Surprisingly, the phenomenon has never been clearly isolated for analysis in its specificity. Thanks to Crispin Wright for directing my attention to it (see his Wright [2010] for some discussion).

and so to be accomplished only throughout (what one would intuitively consider to be) *the whole temporal stretch* of the change, or, at most, only in *considerably large* subintervals thereof. In other words, it seems that the change cannot be *located* in any *considerably smaller* subinterval of the whole temporal stretch of the change (let alone at any *instant* included in it).[52] Indeed, if this were not the case, the very phrase 'the whole temporal stretch of the change', meant to pick out a quite extended temporal interval, would be a dramatic misnomer, as the real change would ultimately consist in a *sudden jump* from red to orange, with, strictly speaking, no real change before and after that (that is, no real change that affects an object's being red or orange). But, taking a cooling bar of iron to be the object changing from red to orange and δ a sufficiently small subinterval, and assuming that the property of being orange is the property acquired by the iron bar as soon as it loses the property of being red, nothing less than the truth of every instance of 'It is not the case that the iron is red at t but not red at $t+\delta$' seems required to rule out an unwanted sudden jump. And that leads in turn to (N^{red}) (notice that the change need not occur along a *temporal* dimension: it may for example occur along a *spatial* one, for instance when the iron bar changes seamlessly from being red at one end to being orange at the other end).

The foregoing assumes for simplicity that the iron bar's change is from being red to being orange, but an analogous point can obviously be made for any alternative candidate to being orange for being the property acquired by the iron bar as soon as it loses the property of being red. Even if it is (rather implausibly) claimed that there is *no first* such property that is so acquired (because the set of relevant properties is not well-ordered by the x-is-acquired-earlier-than-y relation),[53] the point can simply be restated as concerning the change from possession of the property of being red to lack of this very same property (rather than possession of any other property): the change from being red to not being so would seem no less seamless than the one from being red to being orange.

Notice that rejection of the law of excluded middle would indeed allow for rejection of every instance of 'The iron bar is red at t but not red at $t+\delta$'. One rejects, of each member t of a collection of times, that the iron bar is either red at t or not red at t. Assuming very plausibly that rejection of a disjunction requires rejection of both disjuncts, this commits one to rejecting, of each member t of such collection, that the iron bar is red at t and to rejecting, of each member t of such collection, that the iron bar is not red at t. We can assume that these times are strictly later than the times of which one accepts that the iron bar is red at them and strictly earlier than the times of which one accepts that the iron bar is not red at them, that the length of the stretch constituted by these times is at least as great as any admissible appropriately small enough δ and, ignoring

[52] In terms of Vendler [1957]'s influential taxonomy, changing from red to orange is an *accomplishment* (rather than e.g. an achievement). Some accomplishments do not involve seamless change, as *their start point is not exhaustive with their end point* (so that the whole temporal stretch of the change can be used to go through the relevant intermediate categories; think for example of the accomplishment of drawing a circle from scratch). But a few accomplishments (such as the change from red to orange, see Zardini [2016g]) are such that *there is no intermediate category between their start point and their end point*, and it is precisely such accomplishments that involve seamless change.

[53] Throughout, I occasionally use hyphenated open sentences to denote properties and relations.

higher-order vagueness, that these three kinds of times are exhaustive on the relevant domain. Assuming very plausibly that rejection of either conjunct requires rejection of a conjunction, this commits one to rejecting every instance of 'The iron bar is red at t but not red at $t + \delta$'. Interestingly, the converse entailment (from rejection of every instance of 'The iron bar is red at t but not red at $t + \delta$' to rejection of the law of excluded middle) does not hold, as is witnessed by standard supervaluationist approaches (see e.g. Fine [1975]). Indeed, the target of the following four points is the more general position consisting in rejection of every instance of 'The iron bar is red at t but not red at $t + \delta$' and non-acceptance of (N^{red}).

Arguably, rejection of every instance of 'The iron bar is red at t but not red at $t + \delta$' is not sufficient to capture the intuition of seamless change. To begin with, the rejection may be all fine as far as it goes, but in itself it does not offer any *description* of how the change occurs. It is not an *account* of seamless change weaker than and alternative to the one provided by (N^{red}), for it is no account in the first place—it is merely the rejection of one account.[54,55]

More specifically, the intuition in question is, to repeat, that the change *takes time*, and so that it is not located in any considerably smaller subinterval of the whole temporal stretch of the change, which requires, as it were, a "*positive*" lack of change at any such subinterval: of any such subinterval, it is not just rejected that the change occurs in it—it is positively accepted that it is not the case that the change occurs in it. It is positively accepted that the change is, say, a *matter of minutes rather than instants*. But such "positive" lack of change of the iron bar between t and $t + \delta$ requires that it be not the case that the iron bar is red at t but not red at $t + \delta$.

That simple rejection of every instance of 'The iron bar is red at t but not red at $t+\delta$' is not sufficient to capture the intuition of seamless change can also be seen by reflecting on the fact that it is not sufficient to rule out every unwanted sudden jump *of some sort*. For simple rejection that the iron bar is red at t but not red at $t + \delta$ is consistent with e.g. acceptance that the iron bar is red at t but borderline red at $t + \delta$, where 'borderline' can be understood in such a way as to make 'φ and it is borderline that φ' inconsistent, and so in such a way as to make a process in which a red object becomes borderline red enough of a *change*.[56]

[54]By accepting every instance of 'It is not the case that [the iron bar is definitely red at t but definitely not red at $t+\delta$]', we at least get an account of *something* (at least provided that we know what definiteness is supposed to be so that definite red does not boil down to red). But that something is the esoteric seamless change from definite red to definite non-red, not the exoteric seamless change from red to non-red.

[55]Obviously, that is a problem affecting *any* attempt at putting forth a *view* by basically just rejecting a certain thesis (and, obviously, the problem is not adequately solved by additionally accepting the *psychological* claim that one rejects the thesis or the *normative* claim that one should reject the thesis). In some cases, the attempt is at least in some respects ameliorated by grounding the rejection in question in the acceptance of some claim involving e.g. determinacy, but fn 54 indicates that no such amelioration is possible in the case of seamless change.

[56]One could predictably try to rule out also that jump by adding rejection that the iron bar is red at t but borderline red at $t + \delta$. However, such move obviously only invites a second-order analogue of the point in the text to the effect that such rejection is consistent with acceptance that the iron bar is red at

Taking now into consideration also other typical commitments of this scheme, on it, the quantified claim 'For every t, it is not the case that the iron bar is red at t but not red at $t+\delta$' is inconsistent with other uncontroversial assumptions, and so should be rejected, which would seem to leave very little room for manoeuvre for preserving the intuition of seamless change. Intuitively, it's not that we want to reject both that the change occurs instantaneously and that it doesn't (as, for a borderline case of 'red', we might want to reject both that it is red and that it isn't), for, *intuitively, it's not that it's indeterminate whether the change occurs instantaneously*—intuitively, since it's seamless, it definitely doesn't. An analogous consideration concerns the fact that, on any way of implementing this scheme I know of, for some instance of 'The iron bar is red at t but not red at $t+\delta$', rejection of it will require rejection of its negation as well.[57]

The main thrust of the foregoing considerations is that we think of some processes as crucially *taking time*—milk doesn't go off in a nanosecond, Sampras didn't become a great tennis player in a nanosecond, I haven't learnt English in a nanosecond—and, surprisingly enough, nothing less than the truth of the relevant instances of (N) would seem able to entitle us to this very natural conception. Arguably, the conception is meant to be expressed in some uses of the phrase 'seamless change'. It is essential to stress that not every use of this phrase can be made sense of as simply saying that the object in question is changing *continuously* along a certain dimension, exactly because some such use is meant to convey the *negation of sudden jumps in the exemplification of a property*, jumps that are not ruled out by a simple continuous change along a certain dimension. For example, an object could be continuously increasing in temperature: this does not rule out jumps in the exemplification of the relevant heat-related properties, such as the property of being at most 20°C hot. The simple *continuity* of the function that in this case takes a time to the most specific heat-related property exemplified at it does not ensure any *seamlessness* in the sense under discussion. Nor is it required by it, since seamless changes are typically still seamless when seen on film.

The apparatus of exemplifying a property *to a certain degree* could naturally be thought to capture this stronger sense of 'seamless change'. Unfortunately, no satisfactory answer has ever been provided to the question concerning the relation between exemplifying a property *simpliciter* (which is what this discussion is really about) and

t but borderline borderline red at $t+\delta$. One could then predictably try to go through an infinite series of related rejections. However, in themselves, such rejections are not sufficient to rule out *that there is an* unwanted jump of some sort (whereas they are sufficient, for every unwanted jump of some sort, to rule out *that* jump), and, if one tried to overcome this difficulty by pooling them together in the single rejection that the iron bar is red at t but falls in some sense short of being red at $t+\delta$, one would commit to [accepting 'The iron bar is red at $t+\delta$' if one accepts 'The iron bar is red at t'], thereby making resistance to (N^{red}) futile (see Zardini [2014b] for broadly related considerations in a different context).

[57]The proverbial "alert reader" will have realised that relatives of all these four points apply to the adoption of this rejection-based approach in relation e.g. to the issues considered in subsection 2.1, where my argumentation only appealed to a relative of the last point. I present the whole battery of points here as I think that they are even more intuitive when applied to the phenomenon of seamless change. Indeed, one might argue (though I will not attempt to do this here) that this particular phenomenon is somehow a paradigm for our general conception of *tolerance*, just as a venerable tradition has it that the phenomenon of continuous change is somehow a paradigm for our general conception of *continuity*.

exemplifying it to a certain degree (which is what the apparatus is really about). On the one hand, to say that an object exemplifies a property p *simpliciter* iff it exemplifies p at least to degree δ makes no sense of seamless change, since the change from possession of p to lack of p will then be as instantaneous as the change from exemplifying p to degree δ to not doing so. On the other hand, to say that an object exemplifies a property p *simpliciter* iff it exemplifies p at least to a high degree simply shifts the problem of explaining the seamlessness of the change from possession of p to lack of p to the problem of explaining the seamlessness of the change from possession of the newly introduced property of [exemplifying p at least to a high degree] to lack of this latter property. Nothing has been gained. Even worse, the very same notion of exemplifying a property p to a certain degree has usually been explained in such a way as to be reduced to exemplifying *simpliciter* the property of having a certain value along a certain dimension of comparison relevant for p. This reduction entails that the change from exemplifying p to a certain degree to exemplifying it to a different degree, while possibly continuous, has no more to do with seamless change than the change of value along the relevant dimension of comparison has. Since, as we have seen in the last paragraph, the latter is not by itself particularly conducive to seamlessness, it is hard to see how the former could be so.

Those who reject (N) may still hope to make sense of the idea that at least the *preparation* for a change—if not the change itself—takes time. However, it is rather unclear what this idea consists in and how it is supposed to relate to the phenomenon of seamless change. For one thing, a seamless change doesn't feel at all like a long *crescendo* climaxing with the real change (it doesn't feel at all as though there were more change at the end). Should it be added that seamless change includes not only the antecedent preparation for the change, but also the subsequent *accommodation* of the change, it will suffice to observe that nor does a seamless change feel at all like a long *crescendo* climaxing with the real change and followed by a long *decrescendo* (it doesn't feel at all as though there were more change in the middle).

Moreover, even though, once existing, the process in which the preparation consists is allowed to stretch through time, it must be kept in mind that *the property that is eventually going to be lost in the change is still present throughout the preparation for the change and that all the events participating in the preparation also consist in sudden jumps*. Such a preparation does no better in suggesting a seamless change than the uneven journey of an old-fashioned minute hand from 1.00 pm to 1.29 pm does in suggesting a seamless change from the minute hand's not indicating 1.30 pm to its doing so.

Furthermore, from the perspective of a rejection of (N), the preparation for a change, as a kind of process, cannot come more seamlessly into existence than the change itself can. Still, it would seem that the preparation itself, if there is indeed such a thing, might seamlessly come into existence (as in the example of the iron bar's changing from red to orange), and the postulation, required by the current strategy, of *another preparation for the original preparation's coming into existence* looks dubious.

Finally, while the idea of a preparation for a change makes intuitive sense in the case e.g. of the change consisting in the destruction of Carthage (think of the process of destroying houses, burning ships, deporting people etc.), such idea would seem inapplicable

to at least some other cases of seamless change. To return to our original example, there would not seem to be any similar preparation in the iron bar's changing from red to orange—*no set of events that jointly constitute the change even though each of them is in itself insufficient to do so.*

The same conclusions about seamless change can be reached by approaching the phenomenon from a different perspective. It is uncontroversial that we do *not perceive*[58] the sharp boundaries required by what (N) negates. Of course, it is not in general the case that one's not perceiving that P entails one's perceiving that it is not the case that P. For example, being away from home, I may not perceive that my dog is at home without thereby perceiving that he is not. Yet, a case can be made that, in the case of our perception of the sharp boundaries required by what (N) negates, the situation is at least sometimes underdescribed by simply saying that we do not perceive sharp boundaries (I thus disagree with Wright [2007], p. 440, who thinks that this is all that is warranted by the phenomenology). For, at least in some cases, for every pair of neighbouring objects, we do *perceive* that the boundary of the exemplification of a property does *not* fall between them.

For example, in a well-executed slow *diminuendo* from *forte* to *piano*, we can perceive, of any two neighbouring enough moments, that the boundary between the orchestra's playing *forte* and its not doing so does not fall between them. That one can perceive *this* over and above one's not perceiving that the boundary falls between them seems to be what warrants a favourable aesthetic judgement in the first place. For one would not usually think of one's evidence for such judgement to be, peculiarly enough (as opposed to many other very similar aesthetic judgements, like the judgement that the winds are playing too loud), distinctively *second-order* about one's lack of certain perceptions, which is all the relevant evidence that would be afforded were it the case that a sharp boundary is only not perceived to exist rather than also perceived not to exist. And one would not usually think that the orchestra played simply deftly enough to make it the case that one did not perceive a boundary where, for all one could perceive, there might well have been one—rather, the aesthetic judgement is issued only because one would think that the orchestra played deftly enough not to make at any time a sudden jump from its playing *forte* to its not doing so.

Of course, one can accept all these points about the correct description of our phenomenology and still reject (N), on the grounds that our *experience systematically deceives us in this regard.* On this view, our experience of seamless change would be similar to our experience of geometrically impossible situations (as nicely exemplified e.g. in the works of Escher): even though both real *qua* experiences, what they represent is something that cannot be the case. Such move depends on the availability of independent grounds for thinking that what is represented is impossible (which do indeed exist in the case of experiences of geometrically impossible situations). I doubt that such grounds exist in the case of seamless change (I suppose that the Sorites paradox would be the most plausible candidate, but, as I have said in section 1, I think that a plausible weakening of the logic

[58]Here, I use 'perceive' and its like in a non-factive sense, understanding it as synonymous with lengthier constructions such as 'it looks to one as though'.

exists which makes (N) consistent). In any event, recall that the aim of this paper is not so much that of establishing (N) beyond any reasonable doubt, but to expose the unpalatable commitments incurred by rejecting it—in this case, the unpalatable commitment to a new aspect of systematic illusion in the way we experience the world.

3.2 Appearances

In this final subsection, I wish to trace a fairly specific and unusual argumentative path through the unwieldy jungle of questions concerning phenomenal entities (some of the main works relevant to the issues I will touch on are Goodman [1951]; Armstrong [1968]; Jackson and Pinkerton [1973]; Dummett [1975]; Wright [1975]; Wright [1987]; Peacocke [1981]; Linsky [1984]; Travis [1985]; Hardin [1988]; Williamson [1990]; Raffman [2000]; Fara [2001]). I will employ the notion of *phenomenal* identity of x and y with respect to a certain feature f (the idea that x and y *appear* to be the same specific way with respect to f, no matter whether this is believed or not, see Chisholm [1957], pp. 43–53). I will first defend the claim that apparent (i.e. phenomenal) identity requires identity of appearances. I will then contend that apparent identity is preserved across minute enough differences in what determines the relevant feature.[59] I will finally show how these two claims in turn entail certain instances of (N), or something close enough.

Consider the *appearance abstraction* principle (see Zardini [2016b] for a logical treatment of abstraction principles affected by vagueness):

(APP) For every subject s, time t, feature f and all objects x, y, the appearances of x's and y's f for s at t are identical iff [x and y would appear to s at t to be the same specific way with respect to f if presented to s at t].

Before proceeding with my argument, some remarks on (APP) are in order. Firstly, while the occurrence of 'identical' on the left-hand side of (APP) denotes *numerical identity* (among certain exotic objects, appearances), the occurrence of 'same' on their right-hand side denotes *qualitative identity* (among common-or-garden objects, percepts).[60] Secondly, the appearances referred to on the left-hand side are relativised to whichever feature (colour, shape, sound etc.) is referred to on the right-hand side. Thirdly, the counterfactuality of the right-hand side is needed in order to be able to determine appearances to s at t also for objects that are not perceived by s at t. Fourthly, hereafter I will mostly leave implicit the antecedent of the relevant counterfactual conditional (leaving only the modal 'would' to indicate the intended counterfactuality) and the relativisation to subject, time and feature. Fifthly, I will assume as unproblematic the left-to-right direction of (APP).

[59] I understand such differences as being so minute as possibly not to reflect in any difference in the relevant feature itself: for example, I understand a minute enough difference in what determines colour as consisting in a picometrical difference of reflected wavelength.

[60] I will remain neutral here as to whether qualitative identity can itself be reduced to numerical identity (e.g. between properties).

I now turn to the first step of the argument, arguing in favour of the right-to-left direction of (APP). In addition to its extreme intuitive plausibility, this direction would seem essential in fixing the right identity conditions for appearances. For what *explains* the fact that two objects would appear to be the same specific way? It cannot be the fact that their appearances are merely *similar*, because appearances are often similar without determining that the objects they are appearances of would appear to be the same specific way (for example, the appearance of the shape of a collection of 5 grains is very similar to, but not identical with, the appearance of the shape of a collection of 6 grains, but a collection of 5 grains would appear to be of a different shape than a collection of 6 grains). Moreover, even if mere similarity in appearances between two objects could sometimes *negatively* determine that it is not the case that one would appear to be not the same specific way as the other, it is hard to see how it could *positively* determine that one would appear to be the same specific way as the other.[61] It is hard to see, more generally, how any fact falling short of a (numerical) identity could determine that an object would appear to be the same specific way as another object.[62] But, certainly, appearances are the kind of objects related by such identities if anything is.[63]

Indeed, not only does identity of appearances explain what would otherwise seem a (theoretically highly undesirable) primitive appearing of a (qualitative) identity, thus constituting the ultimate explanatory basis of the *explananda* with respect to which appearings of (qualitative) identities are usually appealed to—it is also required for some explanatory work that cannot be carried out by any appearing of a (qualitative) identity. For example, the dog may react in the same specific way upon numerically different, temporally remote but qualitatively very similar calls of his master. There is no rele-

[61] Again, elaborating on a point already made in subsection 3.1, it is not in general true that its being the case that x would not appear to be not the same specific way as y entails its being the case that x would appear to be the same specific way as y—for example, it is not the case that an atom of hydrogen would appear to me now to be of a different shape than an atom of oxygen although it is also not the case that it would appear to me now to be the same shape as an atom of oxygen.

[62] I should stress that the principle appealed to in the text, while very attractive for *appearings*, is doubtlessly rather improbable for other mental states such as *believings*: it is not at all hard to see how some fact falling short of an identity could determine that an object would be *believed* to be the same specific way as another one.

[63] A natural worry against the right-to-left direction of (APP) is that, for example, in a series from red to orange fine-grained enough so that each pair of neighbouring objects would appear to be the same specific way, there is going to be a first object y such that y and the first object z of the series would not appear to be the same specific way. By construction of y, the predecessor x of y is such that it and z would appear to be the same specific way, and it would seem in the spirit of the considerations in the text to assume that such difference between x and y should be explained by a difference in their appearances, despite the fact that, by assumption, x and y would appear to be the same specific way. However, it arguably has to be *vague* which objects in the series are such that they and z would appear to be the same specific way, and, from the point of view of the naive theory of vagueness and of every other theory of vagueness that rejects sharp boundaries, that entails that the worry crucially relies on the untenable assumption that *there is a first object* such that it and z would not appear to be the same specific way. (In contrast, the concept of *belief* might be precise enough so that there is a first object such that it and z would not be believed to be the same specific way (at worst, the subject can be forced to believe in such a precise way as to make that assumption true). But it is rather improbable to assume that the resulting difference should be explained by a difference in appearances (see fn 62).)

vant connection between the two phenomenological states the dog undergoes, let alone a phenomenological state in which it appears to the dog that one object (the first call) sounds the same specific way as another (the second call)—that is, a phenomenological state that would constitute an appearing of a (qualitative) identity. Only the simple, *unapparent* identity of the appearances of the sounds of the two calls is there to explain the identity of the specific reactions of the dog (an analogous example can be given for the inter-subjective rather than inter-temporal case).[64] But, once identity of appearances is admitted in order to explain such inter-temporal and inter-subjective cases, it would become arbitrary not to admit it in order to explain the case where, for some s and t, two objects would appear to s at t to be the same specific way if presented to s at t.

The next step of the argument targets the conclusion that, under epistemically normal conditions,[65] given only minute enough differences between x and y in what determines the relevant feature,[66] an apparent (qualitative) identity between x and y would hold. In addition to its extreme phenomenological compellingness, this conclusion follows from a minimal connection between the epistemic notion of justifiedly believing and the phenomenal notion of appearing. Observe that it should be uncontroversial that, under epistemically normal conditions, minutely enough differing objects presented pairwise would be *justifiedly believed* to be the same specific way, in the sense that the belief that they are the same specific way would be *positively supported* by the available evidence concerning how they appear. To wit, it is not just that the evidence is merely *consistent with* the objects' being the same specific way—rather, the evidence also *points to* the objects' being the same specific way, in the sense of *favouring* the hypothesis that the objects are the same specific way over its negation.

However, if so much is granted, to reject that, under epistemically normal conditions, minutely enough differing objects presented pairwise would appear to be the same specific way would commit one to maintaining that, under epistemically normal conditions, there could be a case of minutely enough differing objects presented pairwise where it is accepted that they are justifiedly believed, on the basis of how they appear, to be the same specific way despite its being rejected that they appear to be so! What could the basis of this justification be, if it is to concern how the objects appear but fall short of being the appearing of a (qualitative) identity? How could the appearing of something less than a (qualitative) identity be the basis for a justified belief that the objects are nothing less than qualitatively identical?[67] It would rather seem that nothing less than the appearing *of a (qualitative) identity* could be an appearing that favours the hypothesis that the objects are the same specific way over its negation, and thus justifies the belief that they

[64] Of course, in such cases the identity of the relevant appearances is not implied by anything like (APP), which only concerns identity of appearances to the *same* subject at the *same* time. I believe that this shows that, important as they are, abstraction principles like (APP) do not exhaust our understanding of appearances.

[65] Recall fn 41 and notice that this qualification gets around the difficulties raised by Fara [2001], pp. 916–920.

[66] Relativisation to the relevant feature will often be left implicit in the following.

[67] Compare: how could the appearing of an object's being orange be the basis for a justified belief that the object is red?

are so. Hence, the uncontroversial fact that, under epistemically normal conditions, if x and y differ minutely enough, the belief that x and y are the same specific way would be justified would seem to imply that, under epistemically normal conditions, if x and y differ minutely enough, x would appear to be the same specific way as y.

Now, drawing the previous two threads together, suppose for concreteness that, under epistemically normal conditions, x and y differ minutely enough in what determines colour. Then, as I have contended in the second step of the argument, x would appear to be the same colour as y, and so, by the right-to-left direction of (APP) I have defended in the first step of the argument, the appearance of the colour of x is identical with the appearance of the colour of y. Therefore, under epistemically normal conditions, if x and y differ minutely enough in what determines colour, the appearance of the colour of x is identical with the appearance of the colour of y—that is, for every subject s and time t constituting an epistemically normal condition, if x and y differ minutely enough in what determines colour, the appearance of the colour of x for s at t is identical with the appearance of the colour of y for s at t. Moreover, given what epistemically normal conditions are, in this conclusion 'x' and 'y' can range over all the elements of the relevant series, which could go for example from a clear case of red to a clear case of orange (notice that this argument nowhere appeals to closure of appearing under logical consequence, and so does not fall prey to the objections levelled by Williamson [1994], pp. 180–184).[68]

Strictly speaking, this conclusion does not yet vindicate (N). It would do so if we could infer 'If x and y differ minutely enough in what determines colour and the appearance of the colour of x is identical with a, the appearance of the colour of y is identical with a' from 'If x and y differ minutely enough in what determines colour, the appearance of the colour of x is identical with the appearance of the colour of y', where the former is in effect (very plausibly) logically equivalent with ($\mathrm{N}^{\textit{is identical with } a}$). The crucial inference that is needed is a variation of the principle of *transitivity of identity*, which, while arguably valid, is invalid in many logics for vagueness (however, unsurprisingly, it is valid in the main logics proposed in Zardini [2016b]). Be that as it may, even if it should implausibly fall short of entailing ($\mathrm{N}^{\textit{is identical with } a}$), the conclusion is certainly in keeping with the spirit of the naive theory of vagueness, since it anyways establishes that *appearances do not discriminate between objects that differ minutely enough.*

4 Conclusion

The foregoing arguments lend a very high plausibility to many instances of (N). In particular, they do so by showing the high value of predicates satisfying (N) in our thought about,

[68] Since the converse of this conclusion (namely, that, under epistemically normal conditions, if the appearance of the colour of x is identical with the appearance of the colour of y, x and y differ minutely enough in what determines colour (if they differ at all)) should be uncontroversial and since identity is arguably transitive, this implies, under some natural assumptions, that x-differs-minutely-enough-from-y is after all transitive. That would seem plausible given the *vagueness* of this relation (it would not of course be plausible for a relation that specifies precisely the amount of difference tolerated). In Zardini [2016b], I show how this much transitivity can crucially still fall short of underwriting paradox.

experience of and interaction with the world. According to the naive theory of vagueness, satisfaction of (N) is, roughly, what the nature of the vagueness of a predicate is, and so, according to the naive theory, the high value so achieved—namely, the possibility of such thoughts, experiences and interactions—is what the point of vague predicates is. We have been exploring in some detail some of the different grounds of this value (the need for classifications that are interesting, flexible and easy; the existence of experiences of slow changes and matching appearances)—these are then what, according to the naive theory, the sources of vagueness are. These sources, we have seen, are rooted in *fundamental facts about our cognition and agency in the world* (like the coarseness of our discriminations and the need of success despite the limitations on our circumstances), which would seem to underlie a *general human hubris of (often) treating something (small) as nothing*. The foregoing arguments thus lend a very high plausibility to the naive theory itself. By accepting (N), it is the only theory that, anchoring vagueness in those fundamental facts, gives it *depth* instead of reducing it in the end—as it happens in all the other theories—to a rather uninteresting nuisance deriving in some way or other from our (or reality's) failure to establish sharp boundaries for our predicates. While on these theories vagueness arises because there is no point in having sharp boundaries, on the naive theory it arises because there is a point (indeed, many) in not having any.

References

William Alston. *The Philosophy of Language*. Prentice Hall, Englewood Cliffs, 1964.

David Armstrong. *A Materialist Theory of Mind*. Routledge, London, 1968.

Linda Burns. Something to do with vagueness. *Southern Journal of Philosophy Supplement*, 33:23–47, 1995.

Roderick Chisholm. *Perceiving: A Philosophical Study*. Cornell University Press, Ithaca, 1957.

Michael Dummett. Wang's paradox. *Synthese*, 30:301–324, 1975.

Delia Fara. Shifting sands: An interest-relative theory of vagueness. *Philosophical Topics*, 28:45–81, 2000.

Delia Fara. Phenomenal continua and the sorites. *Mind*, 110:905–935, 2001.

Hartry Field. No fact of the matter. *Australasian Journal of Philosophy*, 81:457–480, 2003.

Kit Fine. Vagueness, truth and logic. *Synthese*, 30:265–300, 1975.

Nelson Goodman. *The Structure of Appearance*. Harvard University Press, Cambridge MA, 1951.

Clyde Hardin. Phenomenal colors and sorites. *Nous*, 22:213–34, 1988.

Frank Jackson and Ray Pinkerton. On an argument against sensory items. *Mind*, 82: 269–272, 1973.

Hans Kamp. Two theories about adjectives. In Edward Keenan, editor, *Formal Semantics of Natural Language*, pages 123–155. Cambridge University Press, Cambridge, 1975.

George Lakoff. Hedges: A study in meaning criteria and the logic of fuzzy concepts. *Journal of Philosophical Logic*, 2:458–508, 1973.

David Lewis. General semantics. *Synthese*, 22:189–229, 1970.

Bernard Linsky. Phenomenal qualities and the identity of indistinguishables. *Synthese*, 59:363–380, 1984.

Vann McGee. *Truth, Vagueness, and Paradox*. Hackett, Indianapolis, 1991.

Peter Mott. Margins for error and the sorites paradox. *The Philosophical Quarterly*, 48: 494–504, 1998.

Dana Nelkin. The lottery paradox, knowledge and rationality. *The Philosophical Review*, 109:373–409, 2000.

Terence Parsons. Assertion, denial, and the liar paradox. *Journal of Philosophical Logic*, 13:137–152, 1984.

Christopher Peacocke. Are vague predicates incoherent? *Synthese*, 46:121–141, 1981.

Graham Priest. *Doubt Truth to Be a Liar*. Oxford University Press, Oxford, 2006.

Diana Raffman. Is perceptual indiscriminability nontransitive? *Philosophical Topics*, 28: 153–175, 2000.

Stephen Read. Logical consequence as truth-preservation. *Logique et Analyse*, 183:479–493, 2003.

Ian Rumfitt. "Yes" and "No". *Mind*, 109:781–823, 2000.

Mark Sainsbury. Tolerating vagueness. *Proceedings of the Aristotelian Society*, 89:33–48, 1989.

Mark Sainsbury. *Paradoxes*. Cambridge University Press, Cambridge, 2nd edition, 1995.

Wesley Salmon. *The Foundations of Scientific Inference*. University of Pittsburgh Press, Pittsburgh, 1967.

Timothy Smiley. Rejection. *Analysis*, 56:1–9, 1996.

Scott Soames. *Understanding Truth*. Oxford University Press, Oxford, 1999.

Roy Sorensen. The ambiguity of vagueness and precision. *Pacific Philosophical Quarterly*, 70:174–183, 1989.

Roy Sorensen. *Vagueness and Contradiction*. Oxford University Press, Oxford, 2001.

Roy Sorensen. Knowledge beyond the margin for error. *Mind*, 116:717–722, 2007.

Paula Sweeney and Elia Zardini. Vagueness and practical interest. In Paul Égré and Nathan Klinedinst, editors, *Vagueness and Language Use*, pages 249–282. Palgrave MacMillan, Basingstoke, 2011.

Jamie Tappenden. The liar and sorites paradoxes: Toward a unified treatment. *The Journal of Philosophy*, 90:551–577, 1993.

Jamie Tappenden. Negation, denial and language change in philosophical logic. In Dov Gabbay and Heinrich Wansing, editors, *What Is Negation?*, pages 261–298. Kluwer, Dordrecht, 1999.

Charles Travis. Vagueness, observation, and sorites. *Mind*, 94:345–366, 1985.

Peter Unger. There are no ordinary things. *Synthese*, 41:117–154, 1979.

Zeno Vendler. Verbs and times. *The Philosophical Review*, 66:143–160, 1957.

Ruth Weintraub. On sharp boundaries for vague terms. *Synthese*, 138:233–245, 2004.

Timothy Williamson. *Identity and Discrimination*. Blackwell, Oxford, 1990.

Timothy Williamson. Inexact knowledge. *Mind*, 101:217–242, 1992.

Timothy Williamson. *Vagueness*. Routledge, London, 1994.

Timothy Williamson. Margin for error: A reply. *The Philosophical Quarterly*, 50:76–81, 2000a.

Timothy Williamson. *Knowledge and Its Limits*. Oxford University Press, Oxford, 2000b.

Timothy Williamson. Knowledge within the margin for error. *Mind*, 116:723–726, 2007.

Crispin Wright. On the coherence of vague predicates. *Synthese*, 30:325–365, 1975.

Crispin Wright. Language mastery and the sorites paradox. In Gareth Evans and John McDowell, editors, *Truth and Meaning*, pages 223–247. Oxford University Press, Oxford, 1976.

Crispin Wright. Further reflections on the sorites paradox. *Philosophical Topics*, 15: 227–290, 1987.

Crispin Wright. "Wang's paradox". In Randall Auxier and Lewis Hahn, editors, *The Philosophy of Michael Dummett*, pages 415–444. Open Court, Chicago, 2007.

Crispin Wright. The illusion of higher-order vagueness. In Richard Dietz and Sebastiano Moruzzi, editors, *Cuts and Clouds: Essays on the Nature and Logic of Vagueness*, pages 523–549. Oxford University Press, Oxford, 2010.

Lofti Zadeh. Fuzzy logic and approximate reasoning. *Synthese*, 30:407–428, 1975.

Elia Zardini. Squeezing and stretching: How vagueness can outrun borderlineness. *Proceedings of the Aristotelian Society*, 106:419–426, 2006a.

Elia Zardini. Higher-order vagueness and paradox: The glory and misery of S4 definiteness. In Jurğis Šķilters, Matti Eklund, Ólafur Jónsson, and Olav Wiegand, editors, *Paradox: Logical, Cognitive and Communicative Aspects*, volume I of *The Baltic International Yearbook of Cognition, Logic and Communication*, pages 203–220. University of Latvia Press, Riga, 2006b.

Elia Zardini. A model of tolerance. *Studia Logica*, 90:337–368, 2008a.

Elia Zardini. *Living on the Slippery Slope. The Nature, Sources and Logic of Vagueness*. PhD thesis, Department of Logic and Metaphysics, University of St Andrews, 2008b.

Elia Zardini. Towards first-order tolerant logics. In Oleg Prozorov, editor, *Philosophy, Mathematics, Linguistics: Aspects of Interaction*, pages 35–38. Russian Academy of Sciences Press, St Petersburg, 2009.

Elia Zardini. Truth without contra(di)ction. *The Review of Symbolic Logic*, 4:498–535, 2011.

Elia Zardini. Luminosity and vagueness. *Dialectica*, 66:375–410, 2012.

Elia Zardini. Higher-order sorites paradox. *Journal of Philosophical Logic*, 42:25–48, 2013.

Elia Zardini. Evans tolerated. In Kensuke Akiba and Ali Abasnezhad, editors, *Vague Objects and Vague Identity*, pages 327–352. Springer, Dordrecht, 2014a.

Elia Zardini. The general missing from the hierarchy. In Fabio Bacchini, Stefano Caputo, and Massimo Dell'Utri, editors, *New Frontiers in Truth*, pages 176–200. Cambridge Scholars Publishing, Cambridge, 2014b.

Elia Zardini. Breaking the chains. Following-from and transitivity. In Colin Caret and Ole Hjortland, editors, *Foundations of Logical Consequence*, pages 221–275. Oxford University Press, Oxford, 2015a.

Elia Zardini. The opacity of truth. *Topoi*, 34:37–54, 2015b.

Elia Zardini. $K \not\subseteq E$. *Philosophy and Phenomenological Research*, 2016a. Forthcoming.

Elia Zardini. First-order tolerant logics. *The Review of Symbolic Logic*, 2016b. Forthcoming.

Elia Zardini. No state *a priori* known to be factive is mental. *Philosophical Issues*, 2016c. Forthcoming.

Elia Zardini. Closed without boundaries. ms, 2016d.

Elia Zardini. Knowledge is not the most general factive mental state. ms, 2016e.

Elia Zardini. Open knowledge of one's inexact knowledge. ms, 2016f.

Elia Zardini. Red and orange. ms, 2016g.