

AperTO - Archivio Istituzionale Open Access dell'Università di Torino

Truth is vague (in a sense)

This is a pre print version of the following article:

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/1962998> since 2025-09-19T15:52:27Z

Published version:

DOI:10.1413/115407

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

Truth is vague, in a sense

Riccardo Bruni* & Lorenzo Rossi†

Abstract

In this paper, we argue that the notion of truth is vague, in a specific sense. We proceed in the following fashion. First, we offer a new, and arguably natural, characterization of vagueness, based on some formal results which led us to identify as explanatory salient a feature which we refer to as «scalability». Intuitively, a property P is scalable if (a) it induces a relation of P -similarity (similarity regarding P) between individuals, and (b) such relation is associated with a comparison according to some scale. Second, we argue that the new characterization is strictly more inclusive than the traditional ones: if a property is vague in some traditional sense, then it remains vague in ours; however, there are scalable properties that do not count as vague in any of the traditional senses. Finally, we argue that characterizing vagueness via scalability enables us to explain why truth and vagueness share so many intuitive similarities, and are affected by paradoxes that are close relatives of one another. In short, characterizing vagueness via scalability enables us to flesh out the sense in which truth is vague. We close by outlining the implications of considering truth to be scalable (and therefore vague, in our sense).

1 Introduction

A paradox is an argument that, starting from apparently true premises and proceeding via apparently uncontroversial steps, derives an absurd conclusion. Paradoxes display many different features, so much so that one might think that they belong to distinct categories of linguistic and inferential phenomena. Yet, several authors have independently observed that some traits are common to several kinds of paradoxes. This has motivated a growing trend in the literature, which aims at «unifying» paradoxes (in a sense that can, one hopes, be made reasonably precise). Regardless of the specifics, a unification of different kinds of paradoxes would be explanatorily powerful – as it would contribute to identifying which assumptions are responsible

*Dipartimento di Lettere e Filosofia (DILEF), Università degli Studi di Firenze, via della Pergola 60, 50134 Firenze, Italy, riccardo.bruni@unifi.it, ORCID: 0000-0003-2695-0058. This work was supported by the Italian Ministry of Education, University and Research through the PRIN 2017 program “The Manifest Image and the Scientific Image” prot. 2017ZNNW7F_004.

†Dipartimento di Filosofia e Scienze dell’Educazione, Università degli Studi di Torino, Palazzo Nuovo, via Sant’Ottavio 20, 10124 Torino, Italy, lo.rossi@unito.it, ORCID: 0000-0002-1932-5484. This work was supported by the Research Grant MSCA Staff Exchanges 2021 (HORIZON-MSCA-2021-SE-01) no. 1010866295 “PLEXUS: Philosophical, Logical, and Experimental Routes to Substructurality”.

for different kinds of paradoxical arguments – and it would help identifying which solutions to the paradoxes are more general and promising.

This paper contributes to such a unification. More specifically, it offers an explanation of what the common conceptual features of *soritical* paradoxes (i.e. concerning *vague notions*) and *semantic* paradoxes (i.e. concerning semantic notions such as truth) consist in. By «conceptual unification», we mean that the concepts involved – semantic and vague notions – are essentially of the same kind. In a slogan: we will argue that truth is a vague concept, in a specific sense. More precisely, we will argue that truth and vague properties belong to a category, what we call *scalable properties*, that admits a non-artificial characterization – that is, it is a conceptually natural class of properties to isolate – which successfully explains the common traits of truth and vagueness, and of their paradoxes.

The paper is structured as follows. §2 introduces semantic and soritical paradoxes, and motivates the main question of this paper («is truth a vague concept?») in the framework of the unification of paradoxes. §3 introduces the notion of «scalable property» which, we argue, constitutes a promising tool to address the question of the vagueness of the notion of truth. §4 applies the notion of scalable property to vague concepts, arguing that it provides an adequate characterization of vagueness. §5 contains our main argument for the claim that truth is scalable and therefore, in this sense, vague. §6 concludes.

2 Paradoxes: «*In varietate unitas*»

Semantic paradoxes are obtained by the assumption that the truth predicate is *naïve*, i.e. that a sentence φ is equivalent to its truth-predication, i.e. « φ is true». While apparently uncontroversial, the naïveté assumption yields a paradox when we consider a sentence such as «this sentence is not true», which immediately turns out to be true just in case it is not true – a contradiction. This is the *Liar Paradox*, the most well-known semantic paradox. Soritical paradoxes, on the other hand, are obtained from the so-called *tolerance principle*, the plausible thesis that vague predicates tolerate small differences in their conditions of applicability. Let P be a vague predicate («tall», «young», «rich» – what have you): according to the tolerance principle, if P applies to an individual s , and an individual t is very similar to s as far as P is concerned (e.g., for «tall», t might be just 1mm shorter than s), then P applies to t as well. The *Sorites Paradox* shows that repeated applications of this reasoning lead us to declare that clearly non- P individuals are, indeed, P . And this is, again, a contradiction.

Both the origin and the structure of the two arguments have no obvious similarities. However, similar strategies have been independently devised for addressing both kinds of paradoxes, notably the same choices of a non-classical logic.¹ In addi-

¹See, amongst many others, F.G. Asenjo, *A calculus of antinomies*, «Notre Dame journal of formal logic», VII, 1966, pp. 103–105, Jc Beall, *Spandrels of Truth*, Oxford University Press, Oxford, 2009, P. Cobreros, P. Egré, D. Ripley, and R. van Rooij, *Tolerant, classical, strict*, «Journal of philosophical logic», XLI, 2012, pp. 347–85, P. Cobreros, P. Egré, D. Ripley, and R. van Rooij, *Reaching transparent truth*, «Mind», CXXII, 2013, pp. 841–866, P. Cobreros, P. Egré, D. Ripley, and R. van Rooij, *Vagueness, truth and permissive consequence*, in *Unifying the philosophy of truth*, ed. by T. Achourioti, H. Galinon, J. Martínez Fernández, and K. Fujimoto, Springer, Dordrecht, 2015, pp. 409–430, H. Field, *A revenge-immune solution to the semantic paradoxes*,

tion, similar phenomena affect the solutions provided to both kinds of paradoxes, as «solving» a paradox leads to the possibility of re-formulating a similar paradoxical reasoning that exploits that very solution. This dynamics takes the name of *revenge paradox* (for truth) and *higher-order vagueness* (for vagueness).²

These facts have been taken by some authors as an indication that semantic and soritical paradoxes are somehow related, and has prompted some of them to suggest that truth might be a vague concept. However these claims have mostly remained rather speculative, so far. In [ANONYMIZED], we developed a framework where semantic and soritical paradoxes can be shown to be structurally similar, to receive the same treatment throughout a variety of logics, and to derive from the same non-logical assumptions. More specifically, we showed that both naïveté and tolerance can be formally derived from a more general principle, which we call *indiscernibility*. However, the existence of a formal framework to «unify» paradoxes leads naturally to a question: are the notions that give rise to semantic and soritical paradoxes, that is truth and vague predicates respectively, also one of a kind? Put differently: is truth a vague concept?

Before addressing this question, let us clarify the nature of our investigation. The analysis we propose is neither descriptive nor prescriptive. It is not descriptive because we are not conducting an empirical study supported by linguistic or psychological data – we do not offer a theory about how people speak, or think. Moreover, ours is not a prescriptive study either: we do not provide a theory of how people should speak or think. Our investigation is a *conceptual analysis*: by looking at the logico-linguistic behaviour of truth and vague notions, and especially at their paradoxical consequences, we draw some conclusions about the nature and features of these concepts.

3 Scalability

According to the paradox unification offered in [ANONYM.], both truth-theoretical naïveté and tolerance are deducible from a principle of indiscernibility (henceforth IND). IND states that, for a given relation of similarity S , if s and t are similar in the

«Journal of philosophical logic», XXXII, 2003, pp. 139–77, H. Field, *Saving Truth from Paradox*, Oxford University Press, Oxford, 2008, H. Field, *Naïve truth and restricted quantification: saving truth a whole lot better*, «Review of symbolic logic», VII, 2014, pp. 147–191, K. Fine, *Vagueness, truth and logic*, «Synthese», XXX, 1975, pp. 265–300, L. Horsten, *The Tarskian Turn. Deflationism and axiomatic truth*, MIT Press, Cambridge (Mass.), 2012, S. Kripke, *Outline of a theory of truth*, «Journal of philosophy», LXXII, 1975, pp. 690–716, G. Priest, *The logic of paradox*, «Journal of philosophical logic», VIII, 1979, pp. 219–241, G. Priest, *Doubt Truth to be a Liar*, Oxford University Press, Oxford, 2006, D. Ripley, *Conservatively extending classical logic with transparent truth*, «Review of symbolic logic», V, 2012, pp. 354–78, N. Tennant, *Cut for core logic*, «Review of symbolic logic», V, 2012, pp. 450–479, N. Tennant, *A new unified account of truth and paradox*, «Mind», CXXIV, 2015, pp. 571–605, A. Weir, *Naïve truth and sophisticated logic*, in *Deflationism and Paradox*, ed. by Jc Beall and B. Armour-Garb, Oxford University Press, Oxford, 2005, pp. 218–249, E. Zardini, *A model of tolerance*, «Studia Logica», XC, 2008, pp. 337–368. The motto that gives the title to this section is due to E.T. Moneta, Nobel Peace Laureate 1907. The above, non-exhaustive list clearly witnesses the *varietas*, when it comes to paradoxes (and their solutions): we now have to show the *unitas*.

²See, e.g., L. Rossi, *Model-theoretic semantics and revenge paradoxes*, «Philosophical Studies», CLXXVI, 2019, pp. 1035–1054, S. Soames, *Higher-order vagueness for partially defined predicates*, in *Liar and Heaps: New Essays on Paradox*, edited by Jc Beall, Oxford University Press, Oxford, 2004, pp. 128–149.

sense of S , everything that can be said of s can be said of t , and *vice versa*. IND is a schematic principle that requires a relation of similarity to be given. The relation between IND and tolerance is immediate: the former is a slight generalization of the latter. The relation between IND and naïveté is less immediate; nevertheless, it can be shown that the relation between (the names of) sentences that holds when one is a truth-predication of the other generates a version of IND that is sufficient to derive the semantic paradoxes, as it yields a form of naïveté. Semantic and soritical paradoxes, therefore, are both *indiscernibility paradoxes*.

But what happens in an indiscernibility paradox more generally? First, some individuals are declared P -similar, for a property P . Then, a contradiction ensues, showing that declaring them to be similar leads to absurdity.³ So, what exactly goes wrong in assuming that s and t are P -similar? In [ANONYM.] , we suggested that this can be explained as an *error of scale*: s and t have been regarded as P -similar because, implicitly, a scale of P -ness was chosen that was too coarse-grained to let their P -differences emerge; choosing a more fine-grained scale would make their differences apparent, and would suggest that s and t are not P -similar after all. Under this interpretation, an indiscernibility paradox simply evidences that the «wrong» scale has been selected.

If this explanation is correct, it has some intriguing consequences. More specifically, if an indiscernibility paradox witnesses an error of scale, then there must be *scales* that are relevant for the applicability of the notions involved (semantic and vague notions, in our case). Therefore, only properties that induce a relation of similarity according to such properties, and that can be associated with a scale of comparison can yield an indiscernibility paradox. Call such properties *scalable*.⁴ More precisely, by *scalability* we mean the possibility of interpreting similarity relations between individuals (under some respect) as reflecting a comparison according to a scale. For example, «tall» is clearly scalable in our sense, as judgements about similarity in heights of individuals are naturally (perhaps implicitly) interpreted by mapping their height onto a length scale, whose unit of measure is, say, centimeter. We mention the possibility of such an interpretation because, as with our example, mapping a relation to a scale might only be implicit. It is not just the implicit reference to a scale that makes the property scalable though. The scale is not just implicitly referred to, but it can also be «acted upon». Let us clarify this aspect with an example.

Pediatricians used to use the word «term» to indicate that a child was born between the 37th and the 42nd week from conception. The property of being term is scalable since it is based on a measurement on a time scale, whose unit is a full week, which also provides us with the related relation of similarity.⁵ However, as it turned

³We disregard the possible objection that there are various ways to alter the background logic and block the paradoxical derivations since, as shown in [ANONYM.] , *revenge* and *higher-order vagueness* paradoxes reveal that arguments to triviality that are essentially similar to the original semantic and soritical paradoxes can be reproduced in non-classical logics as well.

⁴Our talk of «properties» is devoid of metaphysical implications: by «properties» we simply mean the content of predicates (in the logician's sense), in a way that is compatible with several (if not all the) specific explications of what such content exactly consists in.

⁵Due to the specific character of the property, the relation of similarity it induces is of a peculiar kind, as it admits comparison only between individuals that are term, i.e. that possess the property itself. In particular, the comparison between s and t related to such a similarity relation does not correspond to the possibility of using comparatives at the level of the language, as it makes no sense, of any two individuals s

out, there could be medically relevant differences (e.g. depending on which treatment might be required) between a term child born, say, 37 weeks and 3 days after conception, and a child born during the 40th week. Therefore, the original classification was further refined by introducing the properties «early term» (between the 37th week and 0 days and the 38th week and 6 days), and «full term» (between the 39th week and 0 days of gestation and the 40th week and 6 days).⁶ What happened is clear: the original scale was «zoomed-in», by changing its unit from weeks to weeks and days, in order to use it in more fine-grained measurements. Of course, a scale can also be «zoomed-out», e.g. for the uses where the word «term» is appropriate, and a more fine-grained scale is not required.

Generalizing a little, we could say that the scale associated with a scalable property P determines whether P applies in each given case. There is a «section» of it that determines the cases to which P applies (in our example, the section that goes from week 37 to week 42), and sections that determine the cases to which P does not apply (the part of the time scale before week 37 and after week 42). Of course, there are differences between scalable properties. For instance, contrast the property «term» with «tall». First, while only a «middle» section of the time scale for «term» determines the applicability of the property, «tall» applies presumably from a region of the scale upwards.⁷ Moreover, while the bounds for «being term» are conventionally fixed (based on medical practice), not so for the bounds of «tall». Another important difference among scalable properties revolves around the behavior of the extremes of the relevant section of the associated scales. Those associated with «term» are sharp: 37 week minus one millisecond is still less than 37 week. Things change for «tall» and other scalable properties: while a judgment about tallness seems to require (if implicitly) a comparison between different sections of the scale, the borders of tallness are famously non-sharp. Finally, note that our notion of scalability applies not just to one-dimensional properties like «tall», which vary along just one dimension (in our case, height), but also to multi-dimensional ones like «intelligent», which vary along and are attributed on the basis of several components (such as smartness, capability for logical thinking, and so on). Our notion of scale can have as many dimensions as those along which a property is measured.

So, if the relation of «term» and «tall» with their scales is so different, what it is that makes them both scalable? It seems to us that the unifying, and probably more important feature of scalable properties is that the scales they are associated with can be indefinitely acted upon, by changing their units and refining their regions (whether extremal or not). Tallness, like term-hood, refers to a scale that can be acted upon by changing its unit, thereby zooming in or out. The effects of zooming in or out can again be different in the two cases. For instance, changing the length unit from centimeter to, say, the tenth of a millimeter in the case of tallness may help solving

and t , to say that one of them, say s , is «more term» than t (we thank an anonymous referee for prompting us to clarify this). This is precisely the difference between scalable and *gradable* properties. More on that below in this section.

⁶For the record, the full classification also comprises the properties «late term» (between the 41st week and 0 days and the 41st week and 6 days), and «post-term» (after 42nd week and 0 days).

⁷Note that this description does not commit us to any specific view or cluster of views about the semantics of «tall»: our description is compatible with the region of the scale in question being precisely delimited, fuzzily delimited, context-dependently delimited, pluri-valued, and so on.

some borderline cases between tall and non-tall, thereby sharpening the region of the scale where «tall» ceases to apply. However, even in the case of «term», where the edges are left intact, the property itself can be refined via the introduction of a more fine-grained unit of measure – a refinement that can even find its way into the language, with the introduction of suitable adjectives (as in the case of «early term»). Moreover, it is quite obvious that the zooming in and out of scalable properties can be indefinitely repeated, when borderline cases are identified in the regions of the more fine-grained scale, as it happens in typical examples of higher-order vagueness.

In order to better illustrate the notion of scalable property, let us now contrast it with neighboring notions, namely those of *measurable* and *gradable* properties.

Scalable and measurable properties. We could say that a property is *measurable* if it refers to a notion of measure, such as «being born at exactly 37 weeks and 3 days from conception», «being exactly 1.84 meters tall», «weighing exactly 74 kilos», and so on. Now, there is clearly a measure referred to here, as in many scalable properties. If we use the real line to measure tallness, «being exactly 1.84 meters tall» applies to exactly one real number, one point in the real line. There's no acting upon the scale that can make the application of the property more or less ambiguous, nor can the property be further refined: whether we measure individual heights in nanometers or lightyears, «being exactly 1.84 meters tall» applies to all and only the individuals who are exactly 1.84 meters tall. Therefore, measurable and scalable properties do not coincide, since not all measurable properties are scalable.

Scalable and gradable properties. The notions of scalable property and scalability resemble other concepts that have been widely discussed in connection with vagueness, including linguistic notions such as gradability and scalarity. The main difference between these notions and our notion of scalability is that the former are *linguistic* notions, while ours is not. Therefore the tests that are used by linguists to determine phenomena like scalarity – such as whether an adjective is used in a comparison, or whether superlatives can be constructed for it – may simply not apply, or fail to be relevant to what we call «scalable». To see this, consider again the property «being term», which is scalable, but to which none of the usual linguistic tests to determine gradability can be meaningfully applied.⁸

Nevertheless, there are also clear extensional differences between gradable and scalable. Gradable properties come with a natural idea of «grade» or «degree» associated with them, as the properties expressed by adjectives like «tall», «short», «red», and so on. These properties are also, obviously, scalable in our sense, with the scales being provided by the relevant arrangements of degrees. Gradable properties are however just an example of scalable ones, but the latter include properties where no degree is even implicitly referred to, even though a scale can be associated with them. Beside properties like «being term» of which we have said already, properties such as «being a garment» and «being a means of transportation» are cases in point. These properties are similar to prototypical concepts: in the present writers' everyday

⁸For more details on gradability and the relevant linguistic tests to determine it, see H. Burnett, *Gradability in Natural Language: Logical and Grammatical Foundations*, Oxford University Press, Oxford, 2017.

context, a shirt is a prototypical garment but a toga or a kimono are not; similarly, in the same context, a car is a prototypical means of transportation, but a camel or a hot air balloon are not. «Being a garment» and «being a means of transportation» are therefore scalable properties, in that a scale can be associated with them, by way of comparison with the prototype, that represents the distance between prototypical and non-prototypical instances. Note that such scales display the zooming-in and -out phenomena we described before: depending on whether the context is zoomed in or out, a camel might or might not count as a means of transportation. On the other hand, «being a garment» and «being a means of transportation» are clearly not gradable: they do not involve – not even implicitly – a notion of degree along the lines of those employed by linguists to classify a property (or an adjective) as gradable. That is, even if one were to arrange types of garment or means of transportation in classes according to their proximity to the prototypes, it does not seem reasonable to say that an item is a «garment to degree d », no matter how loose the scale one might want to apply. Finally, note that some non-gradable, scalable prototypical concepts might fall under the category of Carnap’s «comparative concepts», i.e. concepts that «serve for the formulation of the result of a comparison in the form of a more-or-less statement without the use of numerical values».⁹ While our characterization of scalability does not require numerical degrees, characterizing an item as a «garment to degree d » (for possibly non-numerical d) strikes us as inappropriate, while characterizing a shirt as a more typical garment (in the writers’ context) than a kimono seems perfectly natural.

These two considerations show that scalable properties are not identical with gradable ones.¹⁰

4 Scalability and vagueness

Having circumscribed the notion of scalability, we now use it to characterize vagueness. In particular, we will argue for the following conditional:

(*) If a property P is vague, then it is scalable.

Our argument will consist in showing that scalability subsumes the two main characterizations of vagueness currently adopted in the literature, i.e. *admittance of borderline cases* and *tolerance*.¹¹ Therefore, we will argue for the following, more precise

⁹R. Carnap, *Logical Foundations of Probability*, Chicago, Chicago University Press, 1950, p. 9.

¹⁰Another notion in the neighborhood of scalability is fuzziness. The main distinction between the two, however, is that the formal explication of fuzziness consists in the development of specific logics, namely *fuzzy logics*, while scalability comes with no such formal development. Informally, the two notions can be seen to be similar, but there are several ways to model scalability that do not employ fuzzy logics.

¹¹We identify tolerant properties and «soritical» properties, i.e. properties that give rise to soritical paradoxes. To see this, assume that P is soritical. Therefore, there are (s.1) a condition C_P (think: being made of 1.000.000 grains of sand) such that, if s satisfies it, then s is P (think: being a heap); (s.2) a «positive gradient», i.e., a measure γ_P different from 0 (think: 1 grain of sand), such that, if the P -distance between s and t is γ_G (or less), and s satisfies C_P , then t satisfies C_P as well; (s.3) a soritical series S_P , i.e. a series of individuals s_0, \dots, s_n such that (s.3.1) s_0 satisfies C_P , (s.3.2) s_n does not satisfy C_P , (s.3.3) for every k in between 0 and $n - 1$, the P -distance between s_k and s_{k+1} is γ_P . It is clear that if P is soritical, the relation of P -similarity required for tolerance can be extracted from the measure γ_P , and identifying the condition C_P with «being P » immediately provides a version of the tolerance principle. Obtaining soritical properties from tolerant ones, on the other hand, is immediate.

version of (*):

(**) If a property P admits borderline cases or is tolerant, then it is scalable.

If admittance of borderline cases and tolerance are sufficiently good characterizations of vagueness (as it seems), then (**) entails (*).¹² As our arguments show, scalability offers a natural way to incorporate the two main proposed characterizations of vagueness without being «artificial»: it isn't simply their disjunction, but an independent concept that, we argue, does some crucial explanatory work clarifying the main aspects of vague properties and their relation to truth.

Properties that admit borderline cases are scalable. According to a traditional characterization of vagueness, a property P is vague if and only if it admits borderline cases. A *borderline case* for a property P is provided by an individual that is neither clearly P nor clearly not- P . We will argue that if P admits borderline cases then it is scalable. In order to further clarify the relations between scalability and admittance of borderline cases, we will also argue that the converse does not hold.

To see that if P admits borderline cases, then it is scalable, we have to show that we can associate a scale with it, that determines whether P applies to a given individual or not. An obvious possibility here is to define a scale of P -similarity. In order to do that, there are at least three options, depending on how we treat borderline cases. We can use the distinction between clearly P , borderline P , and clearly non- P to define a three-partitioned linear scale, where the clearly P 's are mapped to an extremal point, the borderline P 's to the middle point, and the clearly non- P 's to the other extremal point. Alternatively, we can bundle the borderline P 's together with the clearly P 's or with the clearly non- P 's. Either way, properties that admit borderline cases immediately give rise to scales that, at their simplest, comprise either two or three linearly ordered points – and do so independently of whether the property P is tolerant or not. Once the scale has been fixed, actions on it would precisely produce what characterizes scalable properties: whatever choice has been made about borderline cases for the sake of setting of the scale, zooming it in or out will have the effect that either previously holding applications of P will cease to hold, or new ones will begin to hold instead.

To see that it is not the case that, if P is scalable, then it admits borderline cases, consider again the property of «being term» we mentioned above, with respect to the original scale whose unit of measure are weeks of gestation. As we noticed already, here there is no possible borderline case, since the relevant sections of the time scale are sharp: either an individual s is born at exactly 37th week from conception or after, and s is term, or s is born before the 37th week of gestation, and s is not term. So, some properties are scalable but do not admit borderline cases.¹³

¹²We will also briefly comment on the converse of (**), which arguably does *not* hold, but this is not crucial for our argument.

¹³A further characterization of vagueness comes from supervaluational semantics, i.e. *admitting precisifications*. According to this idea, a property P is vague just in case any interpretation of P admits a precisification, namely a (typically classical) interpretation where every individual is either clearly P or clearly non- P . It is easy to see that this is closely connected with borderline case, as an interpretation of P which has no borderline cases has no non-trivial precisification. For this reason, our arguments above

Tolerant properties are scalable. Let P be a tolerant property. As we have already highlighted, this implies that P comes equipped with a relation of P -similarity \sim_P which is employed to determine any potential cases of P -ness according to tolerance:

$$\text{if } P(s) \text{ and } s \sim_P t, \text{ then } P(t).$$

Tolerance immediately induces a scale, whose features can be read off the extension of the similarity relation \sim_P . Let's consider two examples. First, suppose that \sim_P determines a linear ordering over the domain of \sim_P . Then, any (fine-grained) scale that replicates such an ordering will show that P is scalable. Second, suppose that \sim_P does not determine a linear ordering, e.g. because $r \sim_P s$ and $r \sim_P t$ (for r , s , and t pairwise distinct) but one cannot determine whether r is more similar to s or t ; still, any (coarse-grained) scale replicating this situation by setting the P -distance of r from s to be equal to the one between r and t will also witness the scalability of P . Whether the scale is fine-, or coarse-grained, it should be clear that actions on it would produce the required zooming-in and -out effects typical of scalable properties. A version of this reasoning, essentially, shows that the ordering features, however weak, of the domain of \sim_P immediately entail the scalability of P .

Putting things together, $(**)$ holds, and hence $(*)$ holds too. This, in turn, shows that scalability is a promising characterization of vagueness in at least two respects. First, scalability achieves a greater generality than the one that can be achieved by traditional characterizations via admittance of borderline cases or tolerance alone. Second, as our arguments show, taking scalability to characterize vagueness is not an *ad hoc* move, as relevant scales arise naturally from both borderline cases and the structure of the similarity relation at work in tolerant properties.

The greater generality offered by scalability is coherent with our findings on IND [ANONYM.]. Among the consequences of this characterization, we find that scalable properties include properties that would not be traditionally included amongst the vague ones, but that should be, if their similarities are to be explained. Amongst them, notably, we find semantic notions such as (self-applicable) truth, to which we now turn.

5 Truth as a scalable predicate

Recall that IND applies to a relation of similarity \sim when, if two individuals s and t are \sim -similar, then anything that holds of s also holds of t – more formally, for every open formula $\varphi(x)$ of the language in question, $\varphi(s) \leftrightarrow \varphi(t)$.¹⁴ Now, it is immediate to see that tolerance is just a special case of IND, namely IND restricted to a certain relation of P -similarity and to P -ness: whenever s and t are P -similar, then s is P if and

transfer, *mutatis mutandis*, to precisifications, and so does their conclusion: if a property P admits precisifications, then it is scalable, but the converse does not hold (as witnessed, again, by properties such as «being term»). We thank an anonymous referee for valuable discussion on this point.

¹⁴In non-classical logics, the biconditional might need to be replaced with, say, the possibility to derive $\varphi(t)$ from $\varphi(s)$ and *vice versa*. We gloss over such differences.

only if t is P . What is more surprising is the observation that, beside tolerance, truth-theoretic naïveté – the statement that every sentence φ is equivalent to its own truth predication – is also derivable from IND. Unlike tolerance, naïveté shows no obvious connections with IND. However, a relation of truth-similarity \sim_T and a property P_T can be isolated that provides the required instance of IND. More specifically, say that two terms s and t are truth-similar just in case they code two sentences such that one is the truth-predication of the other. Slightly more formally, $s \sim_T t$ holds if and only if either $s = \ulcorner \psi \urcorner$ and $t = \ulcorner \text{Tr}(\ulcorner \psi \urcorner) \urcorner$, or $t = \ulcorner \psi \urcorner$ and $s = \ulcorner \text{Tr}(\ulcorner \psi \urcorner) \urcorner$ for some formula ψ of our object-language, where $\ulcorner \psi \urcorner$ is the closed term that denotes the code of ψ . The definition of the property P_T is slightly more complex, and requires a definition in a fragment of second-order logic (Δ_1^1 -CA). With truth-similarity and the property P_T at hand, it is possible to show that every instance of naïveté follows from the instance of IND formulated with respect to \sim_T and P_T (in a relatively weak second-order theory).¹⁵

The main consequence of this result is that truth is a scalable property, since it admits a scale derived from the possibility of declaring sentences truth-similar in the sense of \sim_T . This *prima facie* puzzling conclusion is strengthened by the consideration that the instances of indiscernibility that lead to naïveté offer an explanation of the semantic paradoxes (as they do with the paradoxes of scalable properties more generally). If naïveté is just (a consequence of) an instance of IND, the contradictions that derive from naïveté arguably depend on the assumption that a sentence is indiscernible from its truth predication. If this yields triviality, then, it seems natural to conclude that a sentence and its own truth predication should not be regarded as indiscernible, and therefore equivalent. This is the «scale calibration error» that explains the truth-theoretic paradoxes: the scale associated with truth predications has not been set correctly, i.e., in a way that allows the distance between a sentence and its own truth predication to make a difference in inferences.

The diagnosis of the semantic paradoxes, then, is not different from the diagnosis of the paradoxes of any other scalable property. If a property is scalable, then there is a scale, and a scale can always be zoomed in or out, calibrating it differently. Some calibrations might well be unproblematic, as they support instances of IND that do not give rise to paradox. But, generally, the paradoxes reveal the calibrations which are problematic: they indicate that the chosen calibration rendered the scale too coarse in a situation where a more fine-grained one was required.

One might object to this interpretation of our results along the following lines. Our characterization of vague properties as scalable suggests, as per our own description (§3), that any property can be associated with (infinitely) many scales, of different levels of granularity. So, whenever we select a more fine-grained scale in order to differentiate between individuals with respect to a given property P , indiscernibility can kick in again, within the finer-grained distinction, and might yield new versions of the same paradoxes. This means, the objection concludes, that our diagnosis, even if correct, does not conclusively solve the paradoxes.

Now, this very argument, we believe, can actually be used to support our analysis rather than to criticize it. That paradoxes structurally similar to semantic and vagueness-theoretic paradoxes «re-emerge» in any given solution to them, targeting

¹⁵[ANONYM.], Proposition 5.1.

that very solution, is the well-known dynamics of *revenge* and *higher-order vagueness* paradoxes. But, then, this means that our analysis of both kinds of paradoxes as calibration errors, supported by the role of indiscernibility in both kinds of paradoxical arguments, is not just able to account for «standard» semantic and soritical paradoxes, but for their revenge and higher-order versions as well.

This is not to say that our analysis has no critical aspects. The most evident difficulty, for us, is answering the following question: if truth is scalable, then what scale is associated with it? Answering this question brings us into the philosophy of truth and, more specifically, deep into inflationist territory.¹⁶ For deflationists, in fact, inferring « φ is true» from φ is a quasi-logical step that requires no substantial justification nor substantial logico-mathematical resources, and is justified purely by the inferential role of the truth predicate. Therefore, deflationists would reject postulating any semantically relevant difference between φ and « φ is true» – let alone measuring such a difference. But if the inflationists (and we) are right, this is not the case. What does this difference consists in, then, and how is it measured?

Here, the formal analysis of truth comes to the rescue. In a standard, Tarskian approach, the property «being true» for the sentences of a language \mathcal{L} cannot be defined within \mathcal{L} itself. The extension of a truth predicate for a theory T formulated in a language \mathcal{L} , therefore, can only be defined in a language that properly extends \mathcal{L} , and in a theory that is stronger than T . This is because, the Tarskian definition of truth requires a logico-mathematical apparatus that, for any (consistent) theory T formulated in \mathcal{L} , exceeds the deductive resources of T .¹⁷ Therefore, truth predications require a semantic ascent, and their «distance», as it were, from truth-free statements can be measured in terms of the *computational complexity* that is required to construct the very truth predicate they employ.

Summing up: truth is a scalable property, in that it gives rise to a relation of truth-similarity which can be associated with a scale. The relevant scale, we suggest, is provided by the computational complexity of the definition of (an extension for) the truth predicate. Since we know (from Tarski's Theorem) that defining an extension for a truth predicate for any theory T requires computational resources not available in T , interpreting truth-predications requires a more complex theory T' . Measuring the complexity jump required for the definition of truth provides the required scale.

This conclusion has a final, noteworthy consequence: it tells against naïveté. If truth is scalable, then appreciating this holds the key to solving the paradoxes, and this requires relinquishing naïveté. However, and crucially, this does not mean relinquishing the idea that a sentence is equivalent to its truth predication *tout court*: it simply means that such an equivalence needs to take into account a complexity jump (and many theories can formally model this).¹⁸

¹⁶For inflationism, deflationism, and more on the philosophy of truth, see, e.g., Jc Beall and M. Glanzberg, *Where the paths meet: remarks on truth and paradox*, «Midwest studies in philosophy», XXXII, 2008, pp. 169–198.

¹⁷To exemplify, if the object theory is arithmetical, i.e., it is a theory whose axioms comprise assumptions about the basic properties of numbers, then the principles required to define truth for its theorems are set-theoretical in nature and inexpressible by purely (first-order) arithmetical means.

¹⁸One such theory is provided in ANONYM. 2.

6 Conclusion

In this paper, we offered a new characterization of vague properties, based on some novel results which have led us to identify an explanatory salient characteristics of a large class of properties, which we refer to as «scalability». We have shown that the new characterization is more general than the existing ones, in the sense that if a property is vague in some traditional sense – via admittance of borderline cases, or tolerance – then it remains vague in ours (but there are scalable properties that should not be regarded as vague in any of the traditional senses). Focusing on scalability, and offering it as a more comprehensive and yet natural characterization of vagueness, in turn, enables us to explain why truth and vagueness share so many intuitive similarities, and are affected by paradoxes that, as indiscernibility paradoxes, are close relatives of one another. Finally, we outlined the implications of considering truth to be scalable (and therefore vague).

As it is often the case with philosophically guided formal researches, once the conceptual import of the formal results is explored, new technical works is required. The formal relationship we have explored here between indiscernibility, tolerance, and transparency suggests, in particular, the need for a formal framework in which the abstract «scales» at work in scalable properties can be formalized and investigated. This is the task for some future work.