Theories of Truth and Generalized Quantification

Abstract

Generalized quantifiers are quantifiers that express properties of one or more sets. Overwhelming empirical evidence suggests not only that natural language makes use of the expressive power of generalized quantification, but that all natural language quantifiers are generalized quantifiers of a specific type: what are known as *restricted quantifiers*. I discuss how Kripke's theory of truth, and theories that make ancillary use of his logical apparatus, resist the incorporation of such quantifiers. I suggest that this resistance is grounds for concern, noting how some methods for coping with the problem can come into tension with common theoretical aims. Along the way, I raise the question of whether the costs accruing to these theories are worth paying, given that some rival theories can accommodate natural language quantification unproblematically.

On a broadly Fregean understanding, quantifiers express properties of properties.[1] Simplifying, by associating properties with predicate-extensions, quantifiers can be thought of as properties of sets. So the universal quantifier $\forall$ tells us the set of things that ϕ has the property of containing everything. The existential quantifier $\exists$ tells us the set of ϕ has the property of being non-empty. That is, where $\mathcal{L}$ is an interpreted language, $\phi^{\mathcal{L}}$ is the extension of ϕ in $\mathcal{L}$, and M is the domain of $\mathcal{L}$:

$\forall x(\phi x)$ is true in $\mathcal{L}$ if $\phi^{\mathcal{L}} = M$

$\exists x(\phi x)$ is true in $\mathcal{L}$ if $\phi^{\mathcal{L}} \neq \emptyset$

Draft; please don't cite without permission

✉: James.A.R.Shaw@gmail.com

[1] For Frege: higher-order concepts.

Once quantification is thought of in these terms, *generalized quantifiers* are quantifiers that express other properties of sets, or relations between several sets.

There is overwhelming syntactic and semantic evidence that ordinary language quantification makes use of the expressive power afforded by generalized quantification—especially *relational* generalized quantifiers that express a relation between two or more sets. I won't be reviewing this evidence here.[2] Instead, I want to ask: How should our formal theories of truth be influenced by this fact about natural language?

I'll be arguing that many theories of truth—namely, those that rely on Kripkean fixed-point constructions—are poorly positioned to accommodate natural language quantification. I think this puts substantial pressure on such theories, though how this manifests differs from theorist to theorist.

Let me begin with Kripke's formalism. Kripke's theory of truth in part grows out of dissatisfaction with broadly Tarskian theories, on which the apparently univocal word *true* becomes fragmented into a series of truth predicates $true_0$, $true_1$, $true_2$,..., where no truth-predicate (non-vacuously) applies to sentences containing either itself or other truth-predicates with higher subscripts.

That this restriction seemingly flies in the face of ordinary usage and understanding of the truth-predicate was met with some indifference by Tarski, who seemed to take the relevant expressive flexibility of natural language as driving it to inconsistency via paradox-inducing liar sentences like (L).[3]

(L) (L) is false.

The Kripkean view attempts to improve on the Tarskian theory in this respect. The idea, which I will only sketch, is to gradually 'build up' an extension for a single self-applicable truth-predicate in stages, using a series of interpreted languages that allow for 'truth-value gaps' (sentences that are neither true nor false).

In languages accommodating truth-value gaps, predicates P of the language are interpreted not merely by an extension P^+ (the objects to which the predicate truthfully applies) but also an anti-extension P^- (the objects to which the predicate falsely applies).

[2] Classic papers on the topic include Barwise & Cooper (1981), Higginbotham & May (1981), Keenan & Stavi (1986), though the purely logical study of such quantifiers goes back to Mostowski (1957). See Glanzberg (2006) for a recent survey from a philosophical perspective.

[3] Tarski (1936). Though there is considerable controversy over what exactly Tarski meant in saying that natural languages are 'inconsistent'.

When a predicate is applied to an object outside the extension and anti-extension, it receives the third 'gappy' value u.

We begin Kripke's construction with some interpreted language $\mathcal{L}_0$ capable of self-reference, in which a truth-predicate T is given an initial extension/anti-extension assignment $\langle T_0^+, T_0^- \rangle$. This can in principle be anything we like, though a natural starting point is $\langle \emptyset, \emptyset \rangle$—the assignment that treats all applications of the truth-predicate as gappy. We then consider a new interpreted language $\mathcal{L}_1$ which is identical to the original language except that it assigns all true sentences of $\mathcal{L}_0$ to T_1^+—the extension of the truth predicate in $\mathcal{L}_1$—and all false sentences of $\mathcal{L}_0$ to T_1^-—the anti-extension of the the truth predicate in $\mathcal{L}_1$. And we can continue to construct languages by consistently reinterpreting the truth-predicate along this pattern:[4]

$$T_{\alpha+1}^+ = \{\phi \mid \phi \text{ is true in } \mathcal{L}_\alpha\}$$

$$T_{\alpha+1}^- = \{\phi \mid \phi \text{ is false in } \mathcal{L}_\alpha\}$$

If all goes well, then somewhere in the series we arrive at a *fixed-point* language: some $\mathcal{L}_\gamma$ such that if we assign its truths to the extension of the truth-predicate, and its falsehoods to the anti-extension, we just get $\mathcal{L}_\gamma$ back. This will, by construction, contain its own univocal truth-predicate in the following sense:

$$\phi \text{ is true in } \mathcal{L}_\gamma \Leftrightarrow T(\ulcorner \phi \urcorner) \text{ is true in } \mathcal{L}_\gamma$$

The critical issue for our purposes is the condition used to prove that we reach a fixed-point: the *monotonicity* of the process of generating new extensions. That is, we are sure to reach a fixed-point provided the extension and anti-extension of the truth-predicate constantly grow:

$$\forall \alpha : T_\alpha^+ \subseteq T_{\alpha+1}^+ \text{ and } T_\alpha^- \subseteq T_{\alpha+1}^-$$

Whether this occurs depends on two factors: the starting choice of extension/anti-extension pair, and the underlying logical resources of the interpreted languages.

Intriguingly, logical resources that present obstacles to arriving at a fixed point are present in the very first example of Kripke's paper, which is wielded as evidence against the Tarskian theory, namely (1).

[4] Where we extend to transfinite stages with unions of previous extensions (and anti-extensions).

(1) Most of Nixon's assertions about Watergate are false.

As Kripke put it:

> If someone makes such an utterance as (1), he does not attach a subscript, explicit or implicit, to his utterance of 'false', which determines the "level of language" on which he speaks. (Kripke (1975) p.695)

This does seem right, and it arguably ends up preventing Tarski's theory from assigning an appropriate semantics to (1). But, ironically, there are reasons to think Kripke's theory is also incapable of capturing the standard semantics of sentences like (1), but for very different reasons.

(1) makes use of the quantifier *most*, where the customary interpretation of *most* ϕ*s* ψ is that the number of ϕs that ψ outnumber the ϕs that do not ψ (e.g., most dogs bark just in case there are more barking dogs than non-barking dogs).

$$\textit{Most } \phi s\ \psi \text{ is true in } \mathcal{L} \text{ iff } |\phi^{\mathcal{L}} \cap \psi^{\mathcal{L}}| > |\phi^{\mathcal{L}} - \psi^{\mathcal{L}}|$$ [5]

This is a familiar example of a 'fundamentally' binary generalized quantifier—that is, a quantifier whose semantics can't be mimicked with unary quantifiers (and familiar connectives).[6] Since Kripke only discusses first-order logics, strictly speaking he never (even indirectly) explores the semantics of his original motivating sentence.

The problem is that were we to add a quantifier constrained by the above semantics to our base language then, provided that language has the expressive resources for self-reference, we are guaranteed never to reach a fixed-point regardless of our choice of starting extension/anti-extension pair for the truth predicate. The basic problem can be illustrated with the following simple case of semantic paradox:

(a) 1+1=2
(b) Most true sentences in this entire box are above the line.

For any interpreted language $\mathcal{L}_\alpha$, if (b) is not in the extension of the truth-predicate T, then more true sentences are above the line (one) than below (none). That means that (b) will come out true in that language, and should be included in the extension of T for the next language $\mathcal{L}_{\alpha+1}$ in construction.

5 Where $|S|$ denotes the cardinality of the set S.

6 Barwise & Cooper (1981), Keenan & Westerstahl (2011).

By contrast, for any interpreted language $\mathcal{L}_\alpha$, if (b) is in the extension of the truth-predicate T, then exactly half of the true sentences are above the line—so it's no longer the case that most are. That means that (b) will come out false (or perhaps undefined) in that language, and should not figure in the extension of T for the next language $\mathcal{L}_{\alpha+1}$ in the Kripkean construction.

So the value of the quantified sentence here jumps in value from u/f to t and from t to one of u or f in the construction process.[7] As a result, the value never becomes stable. But if this happens for any sentence, we never arrive at a fixed-point language which contains its own truth-predicate in the sense mentioned above.

If one reflects on the issue, there may seem to be nothing special about the binary quantifier *most* in the above example. One might think, for example, that we should be able to recreate problems with a more familiar quantifier like *every*, as below.

(c) 1+1=2
(d) Every true sentence in this entire box is above the line.

In a sense, that's correct. The received syntax and semantics of the English sentence (d) gives it a logical form on which *every* expresses the following binary generalized quantifier.

$$\textit{Every } \phi\ \psi\textit{s} \text{ is true in } \mathcal{L} \text{ iff } \phi^{\mathcal{L}} \subseteq \psi^{\mathcal{L}}$$

This quantifier would also raise trouble for the Kripkean construction, for the same reasons as *most*. This difficulty is circumvented in Kripke's actual constructions by using the familiar logician's rendering of the sentence with unary quantification and a conditional:

$$\forall x(\phi x \to \psi x)$$

The unary quantifier requires every object in the domain to satisfy $\phi x \to \psi x$. And the conditional is in turn given a semantics so that $p \to q$ may receive the value u, provided p gets the value u (as in the Weak and Strong Kleene trivalent schemes). If all this is done, a sentence like (d) won't necessarily bounce back and forth between values, but may remain stably at u, provided an initial interpretation of the truth-predicate places it in a gap.

[7] It is the second jump from t to u or f that creates violations of the monotonicity constraint mentioned before.

The problem, from the perspective of the natural language semanticist, is that this is simply to distort the semantics of the English quantifier *every*. It's true that, unlike *most*, the binary quantifier expressed by *every* does have a reinterpretation in first-order logic. But it only has this in the bivalent (that is, purely true or false) setting. The reduction ceases to be equivalent in the trivalent setting (at least, provided the conditional is given one of the relevant interpretations just mentioned).

This tends to result in certain oddities in the Kripkean system. Familiarly, rendering (2) with a unary quantifier prevents it from being true in the least fixed-point construction (that is gradually built from the empty extension/anti-extension pair).[8]

(2) Every true sentence is true.

Arguably this sentence should not only be a truth, but a *logical* truth—as it would be in any viable system that gave the quantifier in (2) its received natural language semantics.

Though these issues usually go unmentioned in the literature on truth, a rare exploration of their significance can be found in Maudlin (2004) (pp.59–64). Maudlin discusses, in impassioned terms, how logicians use 'logical sleights-of-hand' in indoctrinating new students into the belief that *Every* ϕ ψ*s* involves unrestricted unary quantification by exploiting the relevant equivalence in the binary first-order setting, then often illegitimately generalizing (e.g., to the trivalent case). What underlies our intuitions about the trivial truth of (2), Maudlin claims, is a form of 'restricted' generalized quantification like that I mentioned above.

At this end of this discussion, however, Maudlin notes that his own theory of truth is incapable of accommodating the presence of that quantifier, for reasons very similar to those that lead to problems for fixed-point theories.[9] He concludes: "the [untruth] of [(2)] must be paid as the price for consistency and uniform treatment of the quantifiers." (Maudlin (2004) p.64). He later attempts to mitigate the losses here by allowing that such untrue sentences may still be licensed by *sui generis*, and to some extent arbitrary, pragmatic norms of assertability.

This appeal to arbitrary norms of assertability to resolve problem cases has proven to be among the most controversial features of Maudlin's view,[10] but is too complex

8 For the Weak and Strong Kleene schemes.

9 Maudlin's theory arrives at what is essentially an interpretation of Kripke's minimal fixed-point theory, but via a different, graph-theoretic route. Restricted quantification destabilizes the graph-theoretic semantic dependence relations.

10 See Field (2006) for helpful critical discussion.

an issue to enter into here. For now, I only want to note some simple points about Maudlin's concession. First, the claim that the untruth of (2) is a price that *must* be paid is controversially worded. True, it must be paid *if one wants to maintain Maudlin's theory* (or a representatively a similar one). But, as I'll note shortly, there are at least two theories of truth that can accommodate generalized quantification with minimal, or sometimes no, adjustments. Given this, something more needs to be said by way of justifying this treatment of quantification than: "ordinary quantification conflicts with my formalism." We want to know why this isn't grounds for rejecting that formalism in favor of one that can accommodate the quantifiers.

Maudlin not only overstates the need to jettison such quantifiers, but also understates the costs (in spite of his laudable confrontation of those costs). Maudlin chooses to stress the intuitive grounds for thinking that (2) involves a generalized quantifier, not mentioning the wealth of empirical support for this claim. Perhaps connected with this focus, Maudlin doesn't mention quantifiers other than universals and existentials that also create problems for his (and Kripke's) construction. These include *no*, *many*, *half of*, *several*, *neither*, *both*, the ternary quantifier *more…than* and an indefinite range of cardinal quantifiers like *(exactly) one*, *(exactly) two*, *(exactly) three*, and variants like *at least five*, *at most ten*, *between three and seven*, and so on.

These are all relational quantifiers that share the property of *most* (and *every*, *some*) that create problems for generating fixed-points: they are 'restricted'. Intuitively, this means that the relational quantifier $Q(\phi, \psi)$ is 'only about the ϕs'. More formally, for binary quantifiers $Q(\phi, \psi)$, this is reflected in the quantifier having a property known as *conservativity*:[11]

$$Q(\phi, \psi) \Leftrightarrow Q(\phi, \phi \cap \psi)$$

Most dogs bark just in case most dogs are barking dogs. Few cats bark just in case few cats are barking cats. When a quantifier is conservative in this sense, only the extension of ϕ and the part of ψ contained in that extension—that is, only the ϕs—matter to the truth of the quantified statement.

Intriguingly, empirical evidence suggests that *all natural language quantifiers are*

[11] Suppressing a model-parameter. For ternary quantifiers like *more…than*, conservativity is characterized as follows:

$$Q(\phi, \psi, \xi) \Leftrightarrow Q(\phi, \psi, (\phi \cup \psi) \cap \xi)$$

conservative.[12] Conservativity is thus a candidate to be one of those rare and prized instances of a natural language universal, which is supposed to offer deep insight into our biologically constrained language faculty. Natural language quantification appears, fundamentally, to be restricted quantification. This seems important to bear in mind when considering how to react to the fact that fixed-point (and similar) constructions resist the incorporation of restricted quantification.

Recall that Maudlin proposes to jettison the intuitive reading of (2) in favor of unary quantification, knowing that this imperfectly renders its ordinary form. We now see that if we want to try to incorporate the resources of natural language, we will probably have to do something like this for every natural language quantifier. But, strictly, we can't do this for quantifiers like *most* that had no reduction to unary first-order quantifiers to begin with. So if we aren't to discard these quantifiers entirely, we will have to come up with *ad hoc* distortions of their semantics until we find a related quantifier that can be made to fit in the theory.

I'm not sure what the best way to do this is. But we must at least artificially interfere with the restricted character of the relevant quantifiers. Perhaps we should make the value of the quantifier depend adversely on gappy instances of its 'restrictor' (the first predicate on which the quantifier operates), for example as follows.

> *Most* ϕs ψ is undefined (u) in $\mathcal{L}$ if $\exists o$ such that ϕo is u in $\mathcal{L}$.
>
> Otherwise: *Most* ϕs ψ is true in $\mathcal{L}$ if $|\phi^{\mathcal{L}} \cap \psi^{\mathcal{L}}| > |\phi^{\mathcal{L}} - \psi^{\mathcal{L}}|$ and false if $|\phi^{\mathcal{L}} - \psi^{\mathcal{L}}| \geq |\phi^{\mathcal{L}} \cap \psi^{\mathcal{L}}|$.

Such moves tend to come with intuitive costs. E.g., (III) won't be in the minimal fixed-point.

(I) 1+1=2
(II) 2+2=4
(III) Most true sentences in this box are above the line.

And at any rate, the restriction above isn't sufficient: further distortions are needed to

[12] Barwise & Cooper (1981), Keenan & Stavi (1986). Cf. the more recent discussion in Glanzberg (2006).

avoid problematic violations of monotonicity.[13] I won't explore this any further here.[14] The point I want to stress is that any attempt to accommodate such quantifiers by removing their restrictiveness is a distortion of a pre-existing and completely pervasive feature of natural language quantification.[15]

Just how big a cost is this? It will doubtless depend not only on the nature of the distortion, but the philosopher in question. So far I've focused on Kripke and Maudlin, but the issues here affect a broader range of truth-theorists who make use of fixed-point constructions. These include theorists who endorse an interpretation of some particular fixed-point construction like Soames (1999), theorists that make use of fixed-point constructions as components of a broader view like the contextualist theory of Glanzberg (2004) or the paracomplete theory of Field (2008), and even those theorists who don't accommodate truth-value gaps, but make use of fixed-point constructions in giving consistency proofs, like the dialetheist theories of Priest (2006) or Beall (2009).

All these theorists either owe us modifications (more or less substantial) that incorporate familiar generalized quantification into their theories, or at the very least an explanation of why we aren't entitled to any form of ordinary language quantification in their theories. For different theorists this task will probably play out differently. But the task may increase in urgency for some because of their independent theoretical commitments.

For example, many theorists of truth state one of their aims is to secure 'intuitive' or 'ordinary' reasoning or linguistic use. Field objects to a version of the Strong Kleene

---

[13] E.g., (3) prevents our arrival at a fixed-point on the above semantics irrespective of starting extension/anti-extension pair (and without requiring a form of 'exclusion' negation):

(3) It's not the case that most of the numbered examples in this footnote are true.

[14] I suspect the least problematic way is to give *most* and other quantifiers a kind of supervaluational semantics, which may respect monotonicity. This theory has well-known independent intuitive costs, which I suspect is the reason that, e.g., most authors I discuss below that lean on fixed-points do not use such interpretational schemes.

[15] The situation here might be contrasted with discussions of the failure of fixed-point constructions to accommodate a form of 'exclusion negation' such that $\neg p$ is true when p is gappy. Theorists sometimes soft pedal this conflict by claiming we have no grounds to think exclusion negation is independently coherent, or no grounds to think that it corresponds to expressive resources already present in natural language. Whatever the plausibility of such a response, it cannot be made with respect to the issues about quantification I'm raising.

fixed-point construction alone because it "cripples" ordinary reasoning due to the lack of an adequate conditional.[16] Priest motivates his project as one of supplying a language that can "give its own semantics" because "a natural language...can [do so]."[17] From either the perspective of wanting to safeguard ordinary reasoning, or modeling properties of natural languages, distorting cross-linguistically universal logical and expressive properties of natural language quantifiers seems like it should also count as a serious cost.

Additionally, many authors embrace the view that a (if not *the*) primary function of the truth-predicate is its ability to mimic generalization into sentence position using quantification that ranges only over object position.[18] (For example, using "Everything the pope said is true" to convey $\forall p$(if the pope said that p, then p)). But if a primary point of having a truth predicate is for it to interact with pre-existing forms of quantification, one might regard it as problematic to distort quantificational devices to accommodate that truth-predicate, rather than the other way around. At the very least it seems like we should expect a balance of trade-offs.

I won't explore these issues further here. The point is that the initial costs of jettisoning or warping natural language quantification can be exacerbated by special commitments of various theorists. They should not leave these issues unaddressed.

I've noted that many theorists of truth owe us adjustments or explanations. But I also previewed that some might not. These include the Revision Theory of Gupta & Belnap (1993) and the proceduralist semantics of Gaifman (1992, 2000). Revision Theory can accommodate generalized quantification without any adjustments. Indeed, a key virtue of Revision Theory is that the operation of revision theoretic processes is completely indifferent to the logic of base languages. And proceduralist semantics are free to accommodate restricted quantifiers precisely because they have no reliance on something like fixed-points.[19]

Revision Theory and proceduralist theories come with their own virtues and vices. We shouldn't favor these views over those relying on fixed-point techniques merely because the former neatly accommodate generalized quantifiers. Still, as some theorists

[16] Field (2008) p.73. Cf. Feferman (1984) p.95, or Beall (2009) p.26 on the lack of an adequate conditional in Priest's ***LP***.

[17] Priest (2006), p.70.

[18] E.g.: Field (2008) pp.193, 349, Beall (2009) p.1.

[19] Though Gaifman's particular version of a proceduralist theory requires some minor adjustments to ensure restricted quantifiers fit in smoothly. See [citation omitted] for some discussion.

consider the price of foregoing restricted quantification, it is worth bearing in mind that we have no theoretical grounds to think that our best theory of truth has any such costs to pay.

References

Barwise, Jon & Robin Cooper. 1981. "Generalized Quantifiers and Natural Language." *Linguistics and Philosophy*, vol. 4 (2): 159–219. [2], [4], [8]

Beall, J. C. 2009. *Spandrels of Truth*. Oxford University Press. [9], [10]

Feferman, Solomon. 1984. "Toward Useful Type-Free Theories. I." *Journal of Symbolic Logic*, vol. 49 (1): 75–111. [10]

Field, Hartry. 2006. "Maudlin's Truth and Paradox." *Philosophy and Phenomenological Research*, vol. 73 (3): 713–720. [6]

—. 2008. *Saving Truth From Paradox*. Oxford University Press. [9], [10]

Gaifman, Haim. 1992. "Pointers to Truth." *The Journal of Philosophy*, vol. 89 (5): 223–261. [10]

—. 2000. "Pointers to Propositions." In *Circularity, Definition and Truth*, Anil Gupta & André Chapuis, editors, 79–121. Indian Council Of Philosophical Research. [10]

Glanzberg, Michael. 2004. "A Contextual-Hierarchical Approach to Truth and the Liar Paradox." *Journal of Philosophical Logic*, vol. 33 (1): 27–88. [9]

—. 2006. "Quantifiers." In *The Oxford Handbook of Philosophy of Language*, Ernest Lepore & Barry Smith, editors, 794–821. Oxford University Press. [2], [8]

Gupta, A. & N. Belnap. 1993. *The Revision Theory of Truth*. MIT Press. [10]

Higginbotham, James & Robert May. 1981. "Questions, Quantifiers and Crossing." *Linguistic Review*, vol. 1: 41–80. [2]

Keenan, D. & D. Westerstahl. 2011. "Generalized Quantifiers in Linguistics and Logic." In *Handbook of Logic and Language*, Johan Van Benthem & Alice Ter Meulen, editors, 859fi910. Elsevier. [4]

Keenan, Edward L. & Jonathan Stavi. 1986. "A Semantic Characterization of Natural Language Determiners." *Linguistics and Philosophy*, vol. 9 (3): 253–326. [2], [8]

Kripke, Saul. 1975. "Outline of a Theory of Truth." *The Journal of Philosophy*, vol. 72 (19): 690–716. [4]

Maudlin, Tim. 2004. *Truth and Paradox: Solving the Riddles*. Oxford University Press. [6]

Mostowski, Andrzej. 1957. "On a Generalization of Quantifiers." *Fundamenta Mathematicae*, vol. 44 (2): 12–36. [2]

Priest, Graham. 2006. *In Contradiction: A Study of the Transconsistent*. Oxford University Press. [9], [10]

Soames, Scott. 1999. *Understanding Truth*. Oxford University Press, Oxford. [9]

Tarski, Alfred. 1936. "The Concept of Truth in Formalized Languages." In *Logic, Semantics, Metamathematics*, A. Tarski, editor, 152–278. Oxford University Press. [2]