Domain-sensitivity

Isidora Stojanovic

Abstract   In this paper, I argue that there are good motivations for a relativist account of the domain-sensitivity of quantifier phrases. I will frame the problem as a puzzle involving what looks like a logically valid inference, yet one whose premises are true while the conclusion is false. After discussing some existing accounts, literalist and contextualist, I will present and argue for an account that may be said to be *relativist* in the following sense: (i) a domain of quantification is required for determining truth value, but is idle in determining semantic content, and (ii) the same sentence, as used on one and the same occasion, may receive different truth values relative to different domains.

1 The puzzle of false truths

Suppose that, at a dinner taking place during a philosophy conference, I say to my neighbor:

(1) Everyone is a philosopher.

Suppose furthermore that there are no non-philosophers at that dinner. Then (1) is intuitively true—or, to put it more weakly, (1) has a reading on which it is intuitively true. But now consider the following, uttered right after (1):

(2) Brigitte Bardot is a philosopher.

I. Stojanovic (✉)
Institut Jean-Nicod, CNRS—ENS—EHESS, Paris, France
e-mail: isidora.stojanovic@ens.fr

⊴ Springer

By the rule of universal instantiation, (2) logically follows from (1). Yet (2) is false (on any reading). I will call this an instance of the puzzle of *false truths:* on the assumption that (1) is true and that it logically entails (2), (2) must come out true, yet (2) comes out false.

Formulated differently, the puzzle involves a structural rule that allows adding idle premises to an already valid inference, and ensures, in particular, that 'p, q ⊢ p' is a logically valid inference (since the idle premise 'q' has been added to the Identity axiom 'p ⊢ p'). Consider:

(3) Everyone is a philosopher.
(4) Brigitte Bardot isn't a philosopher.
(5) Therefore, everyone is a philosopher.

Suppose again that (3) is uttered in the context of a dinner party at which no non-philosopher is present. Then (3) will be heard as true, and (4) is also true, but (5), which consists of the same sentence as (3) but is said after (4), will be heard as false.

Although the focus of my paper is on quantifier domains, the puzzle of false truths is not—or, at least, not obviously—restricted to the use of quantifier phrases. In the puzzle of *approximation* (cf. Predelli 2005, pp. 111–16), just as in the case of (1)–(2), we have a truly uttered premise and a falsely uttered conclusion that form a logically valid inference. For, suppose that (6) is said in a context in which what matters is the approximate rather than the exact length of a table that is actually 1,998 mm long:

(6) This table is two meters long.
(7) This table is longer than 1,999 mm.

Then (6) is intuitively true, (7) is false, yet (7) can be logically derived from (6), given that 2m > 1,999 mm.

It is important to distinguish from the outset the apparent failure of logical rules that the above cases seem to demonstrate, from their apparent failure in those cases in which the premise is uttered in one context and the conclusion in a (clearly) different context. The fact that I can truly say "I am sitting now" and then stand up and utter the same sentence falsely is no counter-example to the identity axiom 'p ⊢ p', of which "I am sitting now. Therefore, I am sitting now" is an instance. This is not to say that one cannot *argue* that there is a change in context between (1) and (2), and similarly, between (3)–(4) and (5), and between (6) and (7). In fact, Predelli's own solution to the puzzle of approximation appeals to the idea of a context-shift. But if context-shift there is, it is not an obvious one—not as if (6) were uttered while pointing at one table and (7) while pointing at a different table. By assumption, (2) is uttered in the very same context in which (1) is uttered.

2 Two non-starters: literalism and radical contextualism

The sensitivity of quantifier phrases to contextual domain restriction is a topic on which there is a lot of literature, and the dominant approach that has emerged from it posits some sort of implicit argument for the domain restriction (cf. Westerståhl 1985; Stanley and Williamson 1995; Stanley 2000; Peters and Westerståhl 2006, to

mention only a few). Before I turn to it in the next section, I want to examine two other approaches that have been proposed, and discard them as non-starters.

Literalism will insist that (1) is in fact false, so that (1) ⊢ (2) is an unproblematic instance of the rule of universal instantiation. (Literalism will similarly hold that (3) and (6) are false.) There are two ways of running the literalist strategy. One is to deny that (1), as uttered at the philosophy dinner, is even intuitively true. Such literalists will typically say that on their own intuitions, (1) is false from the outset, although a speaker may *convey* something true by *saying* something false, given the familiar Gricean mechanisms. However, the empirical data presented in the semantic literature appear to show that the majority of speakers do hear utterances such as (1) as true.[1] Another perhaps more plausible way of running the literalist strategy is to acknowledge that (1) is *intuitively* true, but to insist that speakers' intuitions are not a good guide to semantic content, but track some other, pragmatic level—and if so, then the truth value that *semantics* ought to assign to (1) could well be False.[2] The reason why I believe that this approach is a non-starter, is that it violates what I take to be a working assumption in natural language semantics, namely:

Criterion of empirical adequacy:

If the majority of competent speakers are inclined to judge that a given sentence, as used in a given context, is true, then the truth value that the semantic theory predicts for that sentence, with respect to the appropriate assignment of values to the relevant contextual parameters, had better be True.[3]

Radical contextualism will accept that (1) is true and (2) false, but will reject the claim that (1) ⊢ (2) is an instance of a logically valid inference. According to the radical contextualist, the way in which (1) gets assigned value True is not determined by the syntax of (1), the meaning of its expressions, and the semantically represented features of context, but undergoes *unconstrained* pragmatic processes, like free enrichment (cf. Recanati 2004). However, this permeability of semantic value to free pragmatic processes is precisely why this strategy, too, reveals itself as a non-starter. For, it forestalls the very enterprise of systematizing logically valid inferences by means of natural language. For, let A and B be sentences of some natural language, such as English. Then if 'A ⊢ B' should warrant that if A is true, then B must be true (if evaluated with respect to the same parameters), then the pragmatic processes that can be involved in determining the truth value of B had better depend on the pragmatic processes involved in determining A's truth value. But then, those processes are not entirely optional and unconstrained, which is what radical contextualism holds.

---

[1] That is to say, speakers naturally hear such utterances as true, which is not to say that if you subsequently point out to them that, after all, Brigitte Bardot is not a philosopher, they would not change their mind and decide that they might have made a mistake in judging the utterance to be true. (See also Sect. 6 below.)

[2] Literalism of this sort is defended e.g. in Bach 1994; Recanati 2004, or Gillon 2008.

[3] To forestall a possible confusion, this criterion does not imply that semantics actually *predicts* truth values—which would be absurd, as it would imply that semantics tells us which sentences are actually true. What semantics does is provide truth *conditions*—that is, it predicts a truth *value* for a sentence, but only relative to a specification of values for a number of contextual parameters, including the possible-world parameter.

Radical contextualism is, in my opinion, also a non-starter, because it fails to make room for another assumption about formal semantics for natural language that I am taking for granted:

Criterion of conformity to logic:

If $S_1, \ldots, S_n \vdash S$ (that is to say, if S is derivable from $\{S_1, \ldots, S_n\}$ in the deduction system of a chosen logic), and if $[[S_1]]^{M,i} = [[S_2]]^{M,i} = \cdots = [[S_n]]^{M,i} = \text{True}$, then the value of $[[S]]^{M,i}$ had better be True (where '$[[\ldots]]$' stands for the semantic value function, and where M is a structure of interpretation and i a point of evaluation, i.e. a sequence of the relevant parameters of evaluation).

Note that the above criterion crucially involves the notion of logical *validity* $(\vdash)$, which, as it were, falls on the "syntax" side, being a relation among the *sentences* of a formal language, rather than the notion of logical *consequence* $(\models)$, which falls on the "semantics" side.[4] This is important because a radical contextualist could reply that logic deals with contextually determined *contents*, and that the content of e.g. (2) just isn't a logical consequence of the content of (1), or that, similarly, the content of (7) doesn't logically follow from the content of (6). While it does not strike me as absurd to hold that logical *consequence* is a relation among contents, the idea that *validity* is one is extremely implausible, for it forces a view on contents on which they are *syntactic* items that can be used in formulating inference rules. Now, even if we granted this view to radical contextualism, the worry is that the formal language of the resulting logic could simply not be natural language—and this is the main reason why I take radical contextualism to be a non-starter, being incompatible with the very enterprise of formal semantics for natural language (as initiated by Richard Montague).[5]

3 A promising alternative: indexical contextualism

The puzzle of false truths, as it arises in the case of quantifier domain restriction, may be laid down more explicitly as follows:

i. "Everyone is a philosopher. Therefore, Brigitte Bardot is a philosopher" is an instance of the rule of universal instantiation, $\forall xFx \vdash F[x/b]$, where F may be any predicate, and b any individual constant; it is therefore a logically valid inference.

ii. (1) is intuitively true. Hence, if S_1 is the formal representation of the sentence in (1), and if $i_1 = (c_1, s_1)$ where c_1 stands for the context relevant to the interpretation of (1) (presumably, the dinner at the philosophy conference) and s_1 stands

---

[4] An analogous criterion that involved the notion of logical consequence instead would hold trivially, since by definition, $S_1, \ldots, S_n \vdash S$ iff every model that makes $S_1, \ldots, S_n$ true must also make S true.

[5] I have added this paragraph in order to address a worry voiced independently by Philippe De Brabanter and by Mikhail Kissine. Let me stress that the logic-related problems that I see radical contextualism stumble upon are general, the rule of universal instantiation being just one case among many. For example, while A, B $\vdash$ A $\wedge$ B would presumably hold in radical contextualist logic, the logical connective $\wedge$ would not correspond to the English word 'and'—nor, for that matter, to any other word of English or of any other language.

for the circumstances (world, time, etc.) relevant to the determination of (1)'s truth value (presumably, the actual world, the time at which (1) is said, etc.), then, given the criterion of empirical adequacy, we ought to have $[[S_1]]_1^i = \text{True}$.

iii. (2) is intuitively false. Hence, if S_2 is the formal representation of the sentence in (2), and if $i_2 = (c_2, s_2)$ where c_2 stands for the context relevant for the interpretation of (2) and s_2 stands for the circumstances relevant to determining the truth of (2), then we ought to have $[[S_2]]_2^i = \text{False}$.

iv. (1) and (2) are uttered in the same context, and the circumstances relevant to determining their truth values are the same, too. In other words, $i_1 = i_2$.

Claim Given the criterion of conformity to logic, i, ii, iii and iv lead to contradiction.

We have seen that literalism avoids the contradiction by giving up ii, while radical contextualism avoids it by giving up i. *Indexical* contextualism also avoids the contradiction by giving up i, but without giving up the very enterprise of mapping natural language sentences to sentences of a formal language in which logical rules may be expressed. Rather, the idea is that the formal representation associated with the sentence in (1) is not simply '∀xPhilosopher(x)'. To begin with, it has become customary to treat 'everyone' as a restricted quantifier, where the restrictor consists at least of the sortal expressed by the suffix '-one', which, for simplicity, I will take to be the predicate of being human. Then, the formal representation that generalized quantifier theory provides for (1) becomes:

$$[\forall x : \text{Human} (x)] \, \text{Philosopher} (x).$$

However, this will not eliminate the contradiction yet, because the following inference will be valid, yet its premises are true and the conclusion false:

$$[\forall x : \text{Human} (x)] \, \text{Philosopher} (x), \ \text{Human} (b.b.) \ \vdash \text{Philosopher} (b.b.)$$

The gist of the indexical-contextualist move is to suggest that there is a hidden argument in the restricting clause of the quantifier, consisting of a (second-order) variable π that takes predicates as its values. This gives us the following formal representation for (1):

$$[\forall x : \text{Human} (x) \wedge \pi (x)] \, \text{Philosopher} (x).$$

In the context of (1), π will presumably take as its value some complex predicate along the lines of '(to be) at this dinner', and since Brigitte Bardot does not satisfy the predicate, one of the premises in the inference is false, so the inference may remain valid with its conclusion being false.

Indexical contextualism, as its name indicates, handles the context-sensitivity of quantifier expressions in more or less the same way in which it handles the context-sensitivity of indexical and demonstrative pronouns. In both cases, semantic value is supplied by the parameter of the context. In the case of demonstrative pronouns, it will be the pronoun's referent, while in the case of quantifiers, it will be a predicate that

restricts the domain of quantification. Of course, there are important differences. One such difference is in the nature itself of the expression's semantic value: for an indexical pronoun, it is usually an individual, while for a quantifier, it will be a predicate, or a property, or a group of individuals (depending on the details of the analysis). Another difference is that the argument responsible for domain restriction is implicit, i.e. phonetically unrealized, while indexicals are not only explicit, but endowed with a lexically encoded meaning (or with a *character*, in Kaplan's parlance). However, putting such differences aside, one merit of indexical contextualism is that it handles domain-sensitivity by means of existing semantic tools, namely contextual parameters, independently required in order to account for the context-sensitivity of the usual indexicals.

4 Domain-relativity and non-indexical contextualism

The role of supplying semantic values for indexicals and implicit arguments (whether or not we think that there are implicit arguments for quantifier domain restriction) is only one of the two roles that the parameter of context plays in Kaplanian theories. The other role is that of determining the circumstances of evaluation that will, in turn, determine the truth value. By recognizing the two roles of the notion of context, one will easily see another way in which quantifiers may be context-sensitive, a way that does not need to posit any phonetically unrealized arguments in the syntax of sentences containing quantifiers. Consider a standard Kaplanian theory that does not stipulate a unique domain for all possible worlds. In such a theory, the truth value of a sentence containing a quantifier expression, and containing no indexical at all, will still depend on the context, because the context determines the world of evaluation. For example, the sentence "There are over a thousand individuals" will be true if said in, and evaluated at, a world of size over 1.000 (such as the actual world), but false if said in, and evaluated at, a world of size 1.000 or less, even when the two worlds agree on all the facts about the individuals that belong to both worlds' domains.

While it is uncontroversial that in some cases, the domain-sensitivity of quantifier expressions stems from their sensitivity to the world of evaluation (on the assumption that different worlds may specify different domains of individuals), world-sensitivity alone will not yet dispel the puzzle of false truths. For, part of the puzzle is precisely that a sentence such as (1), "Everyone is a philosopher", may, in different contexts, receive different truth values, even when the world of the context, and thereby the world of evaluation, is one and the same, e.g. the actual world.

Although one may try tampering with the possible-world parameter itself, for instance, by using situations rather than worlds, a more natural move would be to include among the circumstances of evaluation a new parameter, on a par with the world and the time of evaluation: a domain of evaluation. Thus (1) will be true when evaluated at domains that do not contain any humans that are not philosophers, false when evaluated at domains containing non-philosophers.[6]

The view that domain-sensitivity can be handled by means of a domain parameter in the circumstances of evaluation, rather than by positing a phonetically unrealized

---

[6] We can leave it open whether (1) should have any truth value when evaluated at an empty domain, and what that truth value would be.

but syntactically represented argument associated with the quantifier, paves the way for a form of relativism—but only if it allows for the idea that the sentence in (1), as used in the context in which it was used, is true if evaluated with respect to a domain that contains only philosophers, and false if evaluated with respect to a domain that also contains a non-philosopher. However, standard Kaplanian theories forestall this possibility, because they endorse the following assumption:[7]

The Contextualist reduction:

The context of use determines the circumstances that, in turn, determine the truth value.

When we pair the contextualist reduction with the view that the circumstances of evaluation specify not only a world and a time but also one or more domains over which quantifiers range, what we obtain is not relativism, but rather, a form of contextualism that, following MacFarlane (2007), came to be called *non-indexical* contextualism. The core feature of non-indexical contextualism is that the context that serves to interpret pronouns and other indexicals also serves to supply the values for the parameters to which truth is relative. Once a sentence is interpreted in a context, it will also have its truth value fixed with respect to the circumstances of that same context, so that the only variations in truth value are variations along the context dimension.

Although Predelli (2005) does not explicitly discuss quantifier domains, the solution that he proposes to the puzzle of approximation is in the spirit of non-indexical contextualism. Part of the idea is that the circumstances of evaluation specify the standards of precision, an idea also found in Lewis (1980). Thus if (6) is evaluated at circumstances in which low standards are at play, it will be true, and if it is evaluated at circumstances with high standards, it is false. But in addition to this, Predelli (2005) endorses the Contextualist reduction, and takes it to be a role of the context to specify the standards of precision. The gist of his proposal is that there is a context-shift from (6) to (7), the idea being that the mere mention of such a precise measure as millimetres sharpens up the standards of precision. When we judged (6) to be true, the relevant context was one where rough standards were at play, but when we judged (7) to be false, the standards became inevitably higher. In the latter context, though, the sentence in (6) is no longer true but false, so that the case described no longer constitutes a potential counter-example to the validity of the inference from (6) to (7).

The Predellian context-shifting strategy, as applied to the puzzle of false truths in the case of quantifiers, tries to reconcile the truth of (1), the falsity of (2), and the logical validity of the inference $S_1 \vdash S_2$ (where S_1 and S_2 are the formal representations for the sentences in (1) and (2)). But, unlike indexical contextualism, non-indexical contextualism does not need to depart from surface syntax. It can continue to take '[∀x: Human(x)]Philosopher(x)' to be an accurate formal representation of (1). What it will say, though, is that (1) and (2) are not really uttered in the same context, hence $i_1 \neq i_2$, which makes it possible to avoid deriving a contradiction from i, ii, iii and iv.

---

[7] The assumption is explicit in Kaplan (1989), and widely endorsed ever since. For criticisms, see e.g. Predelli and Stojanovic (2008).

The idea is that when (2) is uttered, the mere mention of Brigitte Bardot produces a context-shift, since it brings to salience a larger situation, one that besides the people at the dinner also contains Brigitte Bardot. And while we have $[[S_1]]_1^i = $ True, it takes little to see that $[[S_1]]_2^i = $ False. The falsity of (2), then, does not conflict with the truth of (1), because different contexts are responsible for the difference in truth value.

5 Taking stock of the available options

At this stage, we seem to have three promising ways of dispelling the puzzle of false truths in the case of quantifiers:

Indexical contextualism:

reject i: the formal representation of (1) is '[$\forall$x: Human(x) $\wedge$ π(x)] Philosopher(x)', where the value of the second-order variable π is supplied by c_1 (i.e. the value assigned to the context parameter).

Non-indexical contextualism (+ context-shifting):

reject iv: the mere mention of Brigitte Bardot in (2) brings to salience a domain larger than the one relevant in the context of (1), so that the values assigned to the context parameter in determining the truth values of (1) and (2) are different—that is to say, $c_1 \neq c_2$—and, therefore, the circumstances of evaluation can be different, too: $s_1 \neq s_2$, hence $i_1 \neq i_2$.

Relativism:

reject iv: (1) and (2), while interpreted with respect to the same context, are evaluated with respect to different domains. That is to say, we have $i_1 \neq i$ because $s_1 \neq s_2$, although $c_1 = c_2$.

I will present some data that suggest that a sentence containing a quantifier, as used in one and the same context, may be judged true if looked at from its own context, and false if looked at from a different context, and will argue that relativism provides the best analysis of such data. But before we get there, a few clarifications regarding the differences among our options may be helpful.

Indexical and non-indexical contextualism will agree on the truth value to be assigned to a given sentence with respect to a given context, but will disagree on the issue of how this truth value comes to depend upon some particular contextually restricted domain. According to the indexical-contextualist, this dependence is encoded at the level of syntax, while for the non-indexical-contextualist, it only shows up at the stage of semantic evaluation. It is therefore difficult to think of compelling arguments that would adjudicate between the two approaches. One such argument in favor of indexical contextualism is based upon the possibility of *binding* the arguments responsible for domain restriction, as arguably happens in examples such as the following:

(8) At most dinners hosted by Prof. Cheng, everyone is a philosopher.

We naturally understand (8) as true iff most events e such that e is a dinner hosted by Prof. Cheng are also such that everyone at e is a philosopher. This implies that the domain over which the quantifier 'everyone' in (8) ranges is not uniquely determined, but varies as the values ranged over by the quantifier 'most' vary. The indexical-contextualist may handle this by letting the quantifier 'most' bind the variable π that occurs in the restricting clause of 'everyone'. In other words, the formal representation associated with (8) would go along the following lines:

$$[\text{Most } \pi : \pi = \text{'to be at a dinner hosted by Cheng'}]$$
$$([\forall x : \text{Human}(x) \wedge \pi(x)]\text{Philosopher}(x)).$$

This is only a first approximation, since the proposed representation raises some obvious difficulties, such as the need of an explanation of how the quantifier 'most', which seems to quantify over events and is expected to bind first-order variables of the same kind as those bound by 'everyone', becomes a second-order quantifier that can bind predicate variables. Further difficulties are then expected to arise, related to contextual restrictions that now bear on *domains of predicates* over which such second-order quantifiers range, and to their nesting under the scope of yet other quantifiers. Such issues are beyond the scope of the present paper.[8]

What I do want to point out, though, is that sentences such as (8) are not an insuperable obstacle to accounts that handle domain restriction at the level of circumstances of evaluation, as non-indexical contextualism and relativism do. What follows is only a rough outline of how such accounts could handle examples such as (8).

The semantics of quantifiers whose domain restriction appears to depend on the values ranged over by another quantifier of higher scope, of which (8) is an illustration, requires operators that operate on the domain parameter in the circumstances of evaluation, in the same way in which modal operators such as 'it is necessary that' operate on the possible-world parameter, and in which temporal operators such as 'it will always be the case that' operate on the time parameter. The counterpart of the adverbial phrase 'at most dinners hosted by Prof. Cheng' in the formal representation associated with (8) is thus expected to be an operator on the domain parameter. Of course, this operator will be more complex than $\square$ or $\lozenge$ that we know from propositional modal logic: first, because its underlying quantifier is 'most', which cannot be defined in terms of $\forall$ or $\exists$; second, because we will need not a unary but a binary operator, whose truth conditions are defined not only in terms of the embedded clause but also in terms of the restricting clause (in our case, the clause '(to be a) dinner hosted by Prof. Cheng'). What we will want, then, is an operator whose semantics might go as follows:

---

[8] There is a vast literature that, inter alia, addresses precisely those issues. See e.g. Stanley (2000); Stanley and Szabo (2000); Peters and Westerståhl (2006), etc.

$$[[\text{ most } (\phi, \varphi)]]^{c, w, t, d} = \text{ True iff for most d}' \text{such that dRd}', {}^{9}\text{and such that}$$

$$[[\phi]]^{c, w, t, d'} = \text{ True, we have it that } [[\varphi]]^{c, w, t, d'} = \text{ True.}$$

Although this admittedly remains sketchy, the point is that relativism and non-indexical contextualism have sufficient means of handling the so-called binding cases. Moreover, there is no reason to fear that the semantic complexity of the operators deployed must outgrow the complexity to which indexical contextualism is already committed, of which we could get a glimpse through its analysis of the sentence in (8).[10]

Let me now turn to the option that I have called *relativism*. What justifies this label is that a sentence, as used in one and the same context, can still receive different truth values, if evaluated at different circumstances. This is what distinguishes this option from non-indexical contextualism, since the latter stipulates that the context determines the circumstances of evaluation that, in turn, determine the truth value.

The relativist option, as I would like to defend it, is similar to the view defended by John MacFarlane as a response to a variety of issues, including future contingents (cf. MacFarlane 2003), epistemic modals, knowledge ascriptions, and predicates of personal taste. If we were to apply MacFarlane's proposal to the issue of quantifier domain restriction,[11] we would construe the circumstances of evaluation as including a domain parameter, but in addition to this, we would define sentence truth with respect to *two* contexts—a context of *use* and a context of *assessment*. MacFarlane's proposal also involves a principle similar to the Contextualist reduction, except that it is now the context of assessment that, together with the context of utterance, determines the circumstances of evaluation that, in turn, determine the truth value (cf. e.g. MacFarlane 2005, p. 327). In MacFarlane's account, then, the same sentence, as used in one and the same context (of utterance) cannot be assessed from one and the same context (of assessment) and still vary in truth value. In the account that I am proposing, on the other hand, there is no special context-of-assessment parameter, and the question of which values are to be assigned to the various parameters of the circumstances of evaluation in order to obtain a truth value remains entirely open, rather than being settled either by the context in which the sentence was used or by the context from which its truth is being assessed.

---

[9] Just as the semantics of modal and temporal operators requires one or more accessibility relations among worlds and among times, we will need one or more accessibility relations among domains of quantification—that is what the relation R between d and d′ in the definition stands for. While the familiar set-theoretical relations, such as being a subset of or having a non-empty intersection, may be the likely candidates for such accessibility relations, we can leave it open what this relation should specifically be in the case under consideration.

[10] For the general result that modal operators can mimic quantifiers, see e.g. Kuhn (1980). Kuhn shows how there can be a modal logic with the expressive power of first-order logic, so the result might not be straightforwardly applicable to the present case, if indexical contextualism were to opt indeed for using second-order quantifiers to handle the cases of binding. The contention that the binding phenomena can be handled equally well by indexical contextualism, non-indexical contextualism and relativism should therefore be approached with a certain amount of caution.

[11] It should be emphasized that MacFarlane himself has never suggested that his framework should be applied to quantifier domain restriction.

6 Retrospective assessment and domain-sensitivity

It is unlikely that there can be any knock-down arguments in favor of relativism against contextualism. As noted earlier, indexical and non-indexical contextualism will typically agree on the truth value that they assign to a sentence in a context. Relativism, on the other hand, will also agree on that truth value, provided that the circumstances of evaluation at which the sentence gets to be evaluated are the ones corresponding to the context used for interpreting indexicals. Now, when semanticists canvass for ordinary speakers' intuitions on truth value, they are not using abstract triples of the form <sentence, context, circumstance>, or even just pairs <sentence, context>, but rather *utterances*. What this implies is that in most cases, the circumstances relevant to determining the truth value are, by default, those of the context in which the sentence is used. But if we limit ourselves to the question of what is the truth value of a given sentence, as used in a given context, and evaluated at the circumstances of that same context, then relativism and the various forms of contextualism (including even radical contextualism) are going to answer in unison. In other words, if we want to discriminate between relativism and contextualism on semantic grounds, that is, by raising the charge against the one or the other of making inaccurate truth value predictions, then we must look for cases that will allow for evaluating a sentence, as used in some context, at circumstances possibly different from those corresponding to the context in which the sentence is used. Arguments for relativism in areas of language other than domain restriction (e.g. epistemic modals) are usually grounded precisely upon such cases. One of the most compelling phenomena in that respect is *retrospective assessment* (to the extent that there is unity to the phenomenon). Let me introduce it with epistemic modals. Suppose that, asked where Prof. Cheng is, Tarek says:

(9)   She might be in the library.

If Tarek has no evidence to the effect that Cheng is not in the library, then intuitively, he is telling the truth. Now suppose that Tarek learns that Cheng has been in China for the entire month, and says, in reference to his earlier utterance of (9):

(10) That's not true.

(10′) What I then said is false.

Speakers' intuitions are that Tarek is now right in reassessing his earlier utterance of (9) as being false, even if he was equally right in assessing it as being true at the time when he made it. So it seems that the sentence in (9), as used in the context of (9), is true if assessed from the context of (9) but false if assessed from the context of (10) or (10′). Note, *en passant*, that in the case of epistemic modals (unlike domain restriction, as we will see shortly), it does not seem to make much difference whether we substitute the past tense to the present tense in either (10) or (10′): That *was* false. What I then said *was* false.

Whether reassessment data are robust and motivate a relativist approach to epistemic modals is not my immediate concern.[12] Rather, I want to suggest that there exists a similar phenomenon in the case of domain restriction. Consider the following variation on our working example. At a dinner attended by philosophers only, I say:

---

[12] For a defense of (a version) of relativism for epistemic modals, see e.g. Egan et al. (2005).

(11)   Everyone is a philosopher.

Then Brigitte Bardot joins the party, and I say, in reference to my utterance of (11):

(12) That's no longer true.
(12′) What I then said isn't true any more.

While we were inclined to judge the utterance in (11) to be true at the time when it was first made, we are also inclined to think that it is correct to later reassess it as being false, as is done in (12) or (12′). Now, the data are quite subtle and call for clarification. First to be noted is the presence of 'no longer' in (12) and of 'any more' in (12′). While the presence of such phrases is not essential to the argument, it is certainly true that they enhance the reading on which (12) and (12′) come out true. Compare with the following:

(12″)   That's not true. [said in reference to the utterance of (11)].

Speakers who were inclined to judge (11) as being true are either disinclined to judge the reassessment in (12″) as correct, or will find (12″) ambiguous between a (more salient) reading on which it is false, and a (less salient) reading on which it is true, where this true reading is precisely the one made explicit in (12) by adding 'no longer'.[13]

So the claim is not that (12), (12′) or (12″) are true *tout court*, but rather, that they each have a reading on which they are true, and that this reading can be enhanced either by lexical material (such as the introduction of 'no longer') or by pragmatic factors. The question, then, is how to account for the ambiguity and, in particular, for the true reading. I will try to show that relativism (as characterized in the previous two sections) is better equipped to handle the reassessment data than either indexical or non-indexical contextualism. However, we will also see that the sentences in (12), (12′) and (12″) involve constructions whose semantic analysis remains an open issue, which will, in turn, open up ways of assigning value True to the sentences at stake, while remaining within a contextualist approach. In the remainder of this section, I will discuss possible explanations of the ambiguity in (12) and (12′), and in the next section, I will outline the corresponding accounts within the relativist and the contextualist frameworks, then compare them and explain why I believe that the contextualist accounts are less attractive than the relativist account.

People readily say things such as (12) or (12″), yet it is not obvious what such claims about truth really amount to, and what the meaning and the truth conditions of sentences such as "That's not true" are. Part of this indeterminacy might stem from an indeterminacy as to what the word 'that' refers to on such uses. This, in turn, raises the theoretical question of what kind of items can be assigned a truth value. For, 'that' will presumably refer to something that can be said to be true or false. While it is beyond the scope of this paper to address the difficult issue of what the bearers of truth are, let me point to the most obvious ways of disambiguating, or making more explicit, the sentence in (12):

---

[13] There is a small number of speakers who initially judge (11) to be true and who, after it is pointed out to them that Brigitte Bardot isn't a philosopher, judge the reassessment in (12″) as being correct *simpliciter* (rather than ambiguous). However, such speakers are subsequently inclined to display "literalist" intuitions, and to intuitively predict truth values that are insensitive to contextual domain restriction.

i: the utterance itself of (11) used to be true, e.g. at the time of (11), but is no longer true at the time of (12).

ii: the proposition, or the semantic content, expressed by (11) used to be true, e.g. at the time of (11), but is no longer true at the time of (12).

iii: the sentence in (11) used to be true, e.g. when it was interpreted in the context provided by (11), but is no longer true in the context of (12).

I am not claiming that **i**, **ii** and **iii** exhaust the ways in which (12) may be analyzed. Indeed, if one takes 'that' to be a genuine demonstrative, then it should be possible for this demonstrative to stand for any utterance, proposition or sentence (depending on what one allows among the things that may be said to be true or false) that is sufficiently salient in the context of (12).[14] On the other hand, if one thinks that 'that' in (12) is not a referential expression, but rather, a sentential anaphor, then one will likely take (12) to be a shorthand for a more complex sentence such as:

(13)   It is no longer true that everyone is a philosopher.

And (13), in turn, might be expanded as follows:

(13′)   It has been true in the past that everyone is a philosopher, and now it is not true that everyone is a philosopher.

When it comes to interpreting (12′), the range of options appears to be as rich as in the previous case. In general, a report such as "What she said isn't true" may be analyzed, once again, in several ways:

i′:   the utterance that she made isn't true.

ii′:   the proposition that she expressed isn't true.

iii′:   the sentence that she uttered isn't true.

In addition to **i′**, **ii′** and **iii′**, it is again easy to think of further ways of analyzing such speech reports. Thus someone who thinks that "the locution 'what is said' is very far from univocal" (Lewis 1980, p. 97) may suggest that (12′) is true iff some contextually salient utterance, proposition or sentence that stands in the saying relation to the individual at stake is false, where the saying relation itself need not be functional (i.e. there may be many things that are "said" by a given speaker on a given occasion). On the other hand, someone who thinks that the phrase "what so-and-so said" does not refer but behaves as some sort of sentential anaphor may want to expand (12′) into something like:

(14)   I then said that everyone was a philosopher, and it used to be true that everyone was a philosopher, but now it is not true that everyone is a philosopher.

What appears to emerge from the discussion of how to understand (12), (12′) and (12″) is that there are many ways of analyzing such reports, ways that, at a first glance, look equally fine. What is more, I believe that all of the outlined proposals may be

---

[14] For this sort of approach, see Davidson (1968).

made to work, albeit with some stretches and strains.[15] This implies that the weight of arguments coming from considerations about reports such as (12) or (12′) is not going to be overwhelming. But, in a more positive vein, such considerations are still to be taken into account, so that on a final deliberation, when the various pros and cons of the available options are measured and compared, which will be done in the next section, the relativist framework may be expected to have some advantage over the rival accounts.

7 Unstable truth values: a motivation for relativism

In this final section, let me return to our three competitors, viz. indexical and non-indexical contextualism and relativism, and explain why I think that cases such as (12) or (12′) provide some motivation for a relativist framework.

Leaving aside (12′) for the time being, let me begin with an analysis with which indexical contextualism may try to account for the ambiguity of (12) and, in particular, for the true reading. Since our central topic is domain-sensitivity, rather than sensitivity to time, and since the examples at stake crucially involve variations in truth along the temporal dimension, I will take it for granted that the three accounts endorse some form of time-relativity, or, to use philosophers' jargon, some form of temporalism. We may assume that the circumstances of evaluation include a time coordinate, which is the weakest form of temporalism, not believed to be particularly problematic. Remember, though, that this assumption is only for the sake of simplicity, since arguments against a treatment of domain restriction by contextualism should not rest on issues that hinge upon its treatment of time.

Moving straight to what I believe to be the most plausible analysis available to indexical contextualism, recall that the formal representation that it would associate with a sentence such as (11) would roughly go as follows:

$$[\forall x : \text{Human}(x) \wedge \pi(x)] \text{ Philosopher}(x).$$

The context of (11) will then supply a value to π, and plausibly, it will be a predicate like 'attending this dinner party'. The (temporal) proposition that (11), under this assignment of value to π, will express in its context, is one that is true when evaluated at the time of (11), but false when evaluated at the time of (12), given that at that time

there are individuals who satisfy both predicates in the restriction clause but are not philosophers, such as Brigitte Bardot.[16]

Though not implausible *per se*, this analysis stumbles upon the following problem. The context of (11) will supply not one, but two candidate values for the predicate variable π, namely the temporally neutral predicate 'belonging to this dinner party' and the temporally specific predicate '*now* belonging to this dinner party', distinguished from the former in that it contains the temporal indexical 'now'. However, neither the speaker of (11) nor the addressee will have any non-arbitrary reason to prefer either value to the other, since the choice between the two will have no impact on the truth value of (11). But now, note that the true reading of (12), in the scenario envisaged, is only possible if the value assigned to π is the temporally neutral predicate. For, were it the temporally specific predicate, then the only way for (12) to be true is that someone actually belonging to the contextually restricted domain of (11) should cease being a philosopher by the time of (12). If we ignore this case, the ambiguity in (12) and, in particular, its true reading, would then have to depend on the issue of which value the context of (11) *actually* supplied to π. Yet we have seen that there are two equivalent values, so to speak, the choice between which makes no impact on the truth value of (11) and need not be resolved in the context of (11).

In other words, the difficulty that indexical contextualism encounters is that whether the reassessment of a domain-sensitive claim is correct or not will depend on disambiguating what we might call a *referential* ambiguity, having to do with the question of which value the context actually supplied to the predicate variable that is responsible for domain restriction. This suggests that speakers' typical reactions to the question of whether the reassessment in (12) is correct should be to say something like, "Well, let's go and ask the speaker of (11) whether he had the one or the other restriction in mind." Yet this is not the reaction that we find, but even if we did, we would see that both the speaker and his audience are typically undecided between the two, since from their own viewpoint, the choice makes no difference to the truth value.[17]

Turning to non-indexical contextualism, let us first look at how it handles retrospective assessment in domain-insensitive cases. Consider:

(15) Kyoto is the capital of Japan.
(16) That was true a few centuries ago, but it is no longer true.

---

[16] This solution presupposes that 'this dinner party' refers to an object, viz. a dinner party, that is not identified in terms of its members—since, precisely, Bardot is a member of it at the time of (12) but not at the time of (11).

[17] For reasons of space, I must abstrain from considering the various replies that are available to indexical contextualism. Let me discard only one such, which would be to say that as a matter of rule, the supplied value is the temporally neutral one. The reason why I think that this would not work is that there are cases in which the choice does matter. Suppose that at the time of (11) I say "Last night everyone was sick." Under the assignment of the temporally neutral value, what I say is true iff everyone at last night's dinner was sick (last night); under the assignment of the temporally specific value, it is true iff everyone at tonight's dinner was sick last night. This is an ambiguity that has an impact on the truth value, and the speaker and her audience are both aware of it and rely on contextual clues to disambiguate it. So if indexical contextualism were to say that in sentences such as (11), the chosen value is temporally neutral as a matter of rule, it would need to explain why the rule does not apply across the board. Perhaps that can be done, but the point is that indexical contextualism does not have an easy way out of the difficulty pointed out with respect to the reassessment of (11).

Non-indexical contextualism has it that the truth value of (15) is determined with respect to the context of (15), hence given that (15) occurs today, it is False. However, (16) is true because it crucially involves a temporal operator, 'a few centuries ago it was true that', which applies to the proposition associated with (15) and returns value True iff there is time t such that t occurred a few centuries ago and the embedded proposition is true at t.

Now, cases involving domain restriction can be handled on an analogous model:

(17) Everyone is a philosopher.
(18) That's true at this dinner, but it's false at some dinners to which I've been invited.

The truth value of (17) is fixed with respect to the domain restriction operative in the context of (17), hence (by assumption) it is True, but (18) is also true since it contains a domain-shifting operator 'at some dinners such that so-and-so, it is true that', which applies to the proposition associated with (17) and returns value True iff there is a domain d such that so-and-so, and the embedded proposition is true at d.

This gives us both an idea of how non-indexical contextualism might handle the report in (12), repeated below as (19), and of the problems that it will encounter:

(19) That's no longer true. [said in reference to (17)]

The problem is that (19) contains no explicit domain-shifting operator, so that when (19) is evaluated for its truth value, the temporal operator will apply to the proposition associated with (17), giving us True iff that proposition is true at an earlier time and false at the present time, but the proposition associated with (19) will also want to be evaluated at a domain, and it would seem that the domain at which it ought to be evaluated is the domain inherited from (17). That choice, however, will make (19) come out false (assuming that no one ceased being a philosopher from the time of (17) to the time of (19)).

In reply to this problem, a non-indexical contextualist might correctly point out that nothing commits her to the view that the domain relevant to determining the truth value of (19) should be the domain inherited from (17). The gist of non-indexical contextualism is that the context determines the circumstances of evaluation, hence given that (19) is used in a context whose domain includes a non-philosopher (namely Bardot), the proposition that everyone is a philosopher is false when evaluated at the circumstances of (19), which is what we need for (19) to come out true.

Alas, this move, will not quite work. The reason is that if (19) is to come out true, it should also be true, with respect to the circumstances of (19), that the proposition associated (17) has been true in the past. But now, if that proposition is evaluated at the circumstances of (19), whose domain includes Brigitte Bardot, it will come out false. So we are back with the same problem.

It is quite clear from the above discussion that what non-indexical contextualism needs is to treat reports such as (19) as involving some sort of implicit domain-shifting device. One way of doing so would be to take expressions such as 'it used to be true that' to be able to denote not only the usual temporal operators, but also more sophisticated operators that can simultaneously operate on several coordinates, including the domain coordinate, and that we would roughly understand as equivalent to 'there used

to be circumstances such that'.[18] But this solution, too, is not without its problems: it will have to posit ambiguity or, at least, polysemy, in all sorts of expressions that seem perfectly unambiguous (such as 'no longer').

Let us finally see how relativism might analyze the reassessment data in the case of domain restriction. Reconsider (19) (or, equivalently, (12)). Even when said on a particular occasion—e.g. just as Brigitte Bardot joins the philosophers' dinner party—it will not have a fixed truth value, but will only receive one when evaluated at particular circumstances, specifying world, time and domain. If the first conjunct, viz. that it has been true in the past that everyone is a philosopher, is evaluated at the domain restricted to the dinner party as of (17), and if the second conjunct, viz. that it is now not true that everyone is a philosopher, is evaluated at the larger domain of (19), which includes Brigitte Bardot, then the outcome is value True. But if both conjuncts are evaluated at the domain of (19), the outcome will be value False, and similarly, if both conjuncts are evaluated at the domain of (17), the outcome is again False.[19]

Relativism has thus a clear advantage over contextualism in that it can straightforwardly predict the ambiguity in (12), (12′) and (19). While we may be willing to accept that both indexical and non-indexical contextualism could come up with ways of accounting for the ambiguity at stake, we must also recognize that, unlike relativism, such contextualist moves would require substantial amendments to the original account.

8 Concluding remarks

This paper has been concerned with two difficult topics: that of contextual restrictions on domains of quantification, and that of the debate between literalism, various forms of contextualism, and relativism. The complexity of both topics makes it inevitable that certain views might not have received the attention they deserve, and that certain issues have not been treated in as much detail as it would have been desirable otherwise. However, I hope to have brought to light some interesting new aspects of domain-sensitivity. First, I have shown that the problem may be formulated as a puzzle involving (what looks like) a logically valid inference whose premises are true while the conclusion is false, and I have pointed out that the puzzle is not intrinsically linked to domain restriction, but appears to concern broader areas of language. Second, I have argued that domain restriction is of interest to the contextualism/relativism debate: the same type of arguments as are usually brought up in the debate is applicable

---

[18] Another option would be to posit covert domain-shifting operators, though this would deprive non-indexical contextualism from an advantage that it had over indexical contextualism, namely, that it did not need to depart from surface syntax.

[19] Note that relativism allows, in principle, that the first conjunct gets evaluated at the domain of (19) and the second, at the domain of (17); but that will hardly ever happen in practice, and it would be interesting to know whether this theoretically possible evaluation of (19) is forestalled merely as a matter of pragmatics, or in part because of lexically encoded factors. For, the true reading of (19) is obtained by evaluating the conjunct in the past tense at a domain that corresponds to the pertinent moment in the past (viz. the domain of (17)), and the conjunct in the present tense, at a present domain (that of (19)). This is certainly not accidental, but there remains the issue of whether this correspondence results from a pragmatic process, or may be traced to some syntactic or semantic level.

to domain restriction, for it appears possible to correctly reassess as false a claim that, assessed from its original context, was true. However, the data regarding reassessment are very subtle, and involve constructions whose syntax, semantics and pragmatics are still under study. The many ways in which it is possible to analyze sentences such as "That's true", "That's no longer true", "What you said is false", and the like, a brief discussion of which was provided in Sect. 6, suggest that any arguments that rely on the cases of retrospective assessment should be received with caution.[20] The goal of the discussion in Sect. 7 was to explain what kind of motivation relativism may draw from the interaction of sentences such as "That's no longer true" and sentences containing quantifiers. I have argued that it can provide an explanation of what is going on in such cases, without amending or modifying the original framework. Of course, this will not constitute a knock-down argument against the rival forms of contextualism, indexical and non-indexical, both of which may eventually find ways of accounting for the data under consideration. But the fact that the relativist approach works well in those cases brings a new element to the stack of issues that ought to be considered and weighed in choosing a semantic framework. The final thought, then, is that on an overall assessment, relativist semantics is likely to be better off, or, at any rate, not any worse off, than contextualist semantics.

References

Bach, K. (1994). Seemingly semantic intuitions. In J. Campbell et al. (Eds.), *Meaning and truth* (pp. 21–33). New York: Seven Bridges Press.

Davidson, D. (1968). On saying that. In *Inquiries into truth and interpretation* (pp. 93–108). Oxford: Clarendon Press (1984).

Egan, A., Hawthorne, J., & Weatherson, B. (2005). Epistemic modals in context. In G. Preyer & G. Peter (Eds.), *Contextualism in philosophy* (pp. 131–168). Oxford: Oxford University Press.

Gillon, B. (2008). On the semantics/pragmatics distinction. *Synthese, 165*, 373–384.

Kaplan, D. (1989) Demonstratives: an essay on the semantics, logic, metaphysics, and epistemology of demonstratives and other indexicals. In J. Almog, J. Perry, & H. Wettstein (Eds.), *Themes from Kaplan* (pp. 481–563), Oxford University Press.

Kuhn, S. (1980). Quantifiers as modal operators. *Studia Logica, 39*, 145–158.

Lewis, S. (1980). Index, context and content. In S. Kanger & S. Ohman (Eds.), *Philosophy and grammar* (pp. 78–100). Dordrecht: Reidel.

MacFarlane, J. (2003). Future contingents and relative truth. *The Philosophical Quarterly, 53*, 321–336.

MacFarlane, J. (2005). Making sense of relative truth. *Proceedings of the Aristotelian Society, 105*, 321–339.

MacFarlane, J. (2007). Nonindexical contextualism. *Synthese*. doi:10.1007/s11229-007-9286-2.

Peters, S., & Westerståhl, D. (2006). *Quantifiers in language and logic*. Oxford: Clarendon Press.

Predelli, S. (2005). *Contexts: Meaning, truth, and the use of language*. Oxford: Clarendon Press.

Predelli, S., & Stojanovic, I. (2008). Semantic relativism and the logic of indexicals. In M. García-Carpintero & M. Kölbel (Eds.), *Relative truth* (pp. 63–80). Oxford: Oxford University Press.

---

[20] The same remark applies to arguments built upon the cases of third party assessments, in which a claim that is true from the viewpoint of the conversation participants is correctly assessed as false by an outside observer.

Recanati, F. (2004). *Literal meaning*. Cambridge: Cambridge University Press.

Stanley, J. (2000). Context and logical form. *Linguistics and Philosophy, 23*, 391–434.

Stanley, J., & Szabo, Z. (2000). On quantifier domain restriction. *Mind and Language, 15*, 219–261.

Stanley, J., & Williamson, T. (1995). Quantifiers and context-dependence. *Analysis, 55*, 291–295.

Stojanovic, I. (2007). What is said as lexical meaning. *Cadernos de Filosofia, 21*, 7–42.

Stojanovic, I. (2008). *What is said: an inquiry into reference, meaning, and content*. Saarbrücken: VDM Verlag.

Westerståhl, D. (1985). Determiners and context sets. In J. van Benthem & A. ter Meulen (Eds.), *Generalized quantifiers in natural language* (pp. 45–71). Dordrecht: Foris.