1 An original sin?

If we want compliance to be a purely logical notion, the compliant responses to a sentence should not depend on any specific sort of "world knowledge". Rather, they should be defined by referring to *all* possible interpretations all the language.

When I set up the framework for the boundedness counterexample, I considered a particular kind of context, where certain facts are known, namely, what individuals the domain consists of, what is the denotation of each individual constant, and, crucially, what is the denotation of the function symbol $+$. Notice that, from the purely logical point of view, any two constants c and c' are perfectly indistinguishable. This ensures that, if $\varphi(x)$ is a formula where neither c nor c' occurs, then the relation between $\varphi(c)$ and $\varphi(c')$ must be perfectly symmetric: for instance, in terms of entailment, $\varphi(c)$ and $\varphi(c')$ must be either incomparable, or equivalent. Knowledge of the denotation of constants *and* the symbol $+$ destroys this symmetry, allowing the construction of a formula $B(x)$ such that, say, $B(0)$ entails $B(1)$ but not vice versa. This asymmetry can be employed to generate an infinite chain of weaker and weaker formulas, all of the form $B(n)$ for some constant n.

However, it is not at all clear that the notion of compliance should take into account the specific interpretation of the symbol $+$, or any symbol, for that matter. In fact, it is much more natural to think that the compliant responses should be generated depending only on the general logical form of sentences. Notice that, for any natural numer n, the entailment from $B(n)$ to $B(n+1)$ is not logical in nature, in that it is only valid given a specific interpretation of the language. But it is quite natural to think that the notion of compliant responses should be utterly blind to any specific interpretation of the language.

In considering the boundedness case, my idea was the following. The semantics over a fixed discourse structure will be a particular case of a more general semantics where domain and interpretation of constant and function symbols will also be at stake. Now, I thought, if troubles arise within this more particular case, they will also arise in the more general setting. For, given the general semantics, we can set ourselves in a context in which the relevant information about domain, constants, and function symbols is part of the common ground. In this way, we are back to the special case. Then, if we try to compute compliant responses in this context, we will be in trouble. But I did not consider the idea that we might not need to compute alternatives with respect to a specific common ground in the first place.

Of course, what happens in a specific information state is interesting, but it need not interfere with the definition of (logically) compliant responses. Take the boundedness formula: purely logically, we get that $B(0), B(1)$, etc. are the compliant responses. In a particular common ground, they may well entail each other. If this happens, then far from being un-compliant, the more informative among these responses can be said to be *more* compliant. We might also have sentence that come to be "pragmatically" compliant. But all such pragmatic notions can enter the picture at a later stage, and their definition will likely rely on a general notion of logically compliant responses. Let me then try and sketch a general semantics in which this logical notion of compliance may be framed.

2 First-order inquisitive semantics

Let $\mathcal{L}$ be a first-order language. For simplicity, let me assume that this language does not contain any function symbols of positive arity. Let us consider a set $\mathcal{D}$ of *possible* individuals, which we may assume to be infinite.[1] Following Kripke, I will also assume that constants are interpreted rigidly. That is, I assume an interpretation I that maps any constant a to a possible individual $I(a)$.[2]

Definition 2.1 (Possible worlds)**.** A possible world w consists of:

1. a domain $D_w \subseteq \mathcal{D}$;
2. interpretations $I_w(R) \subseteq (D_w)^n$ of any relation symbol R in the language.

For any individual $d \in \mathcal{D}$, let $|d|$ denote the state $\{w \mid d \in D_w\}$ incorporating the information that the element d is in the domain.

Definition 2.2 (States)**.** A *state* is a set of possible worlds.

To such a state, we can associate two sets of individuals. One is the set of individuals which are known to be in the domain, that is, those individuals which belong to the domain of every world $w \in s$. The other is the set of individuals which *might* be in the domain, that is, those individuals which are in the domain of at least one world $w \in s$.

Definition 2.3 (Explicit domain of a state)**.** Let s be a state:

- the set of established elements in s is $D_s^{\cap} = \bigcap_{w \in s} D_w$
- the set of possible elements in s is $D_s^{\cup} = \bigcup_{w \in s} D_w$

Notice that as we enhance our information state, we establish more and more elements as belonging to the domain, and discard more and more elements as not belonging to the domain. That is, we have the following fact.

Proposition 2.4. if $t \subseteq s$, then $D_t^{\cap} \supseteq D_s^{\cap}$ and $D_t^{\cup} \subseteq D_s^{\cup}$.

An assignment will be a map g that assigns to any variable x a value $g(x) \in \mathcal{D}$. The interpretation $[t]_g$ of a term t with respect to the assignment g is defined as: $[t]_g = g(t)$, if t is a variable, and $[t]_g = I(t)$, if t is a constant. If $\vec{t} = \langle t_1, \ldots, t_n \rangle$ then $[\vec{t}]_g = \langle [t_1]_g, \ldots, [t_n]_g \rangle$. Relative to an assignment g, the support relation is defined as follows.

[1] The infinity of $\mathcal{D}$ is merely required to prevent there being validities that arise merely as a consequence of the size of $\mathcal{D}$. The precise infinite cardinality of $\mathcal{D}$ is not important, since the Löwenheim-Skolem theorems ensure that first-order logic is insensitive to differences in infinite cardinalities.

[2] This assumption should be innocent enough in terms of logical repercussions, but, should it turn out otherwise, it can always be removed.

Definition 2.5 (Support).

1. $s \models_g R\vec{t} \iff$ for all $w \in s$, $[\vec{t}]_g \in I_w(R)$[3]
2. $s \models_g \varphi \vee \psi \iff s \models_g \varphi$ or $s \models_g \psi$
3. $s \models_g \varphi \wedge \psi \iff s \models_g \varphi$ and $s \models_g \psi$
4. $s \models_g \varphi \to \psi \iff$ for every $t \subseteq s$, if $t \models_g \varphi$ then $t \models_g \psi$
5. $s \models_g \exists x\varphi \iff$ for some $d \in D^{\cap}_s$, $s \models_{g[x\mapsto d]} \varphi$
6. $s \models_g \forall x\varphi \iff$ for all $t \subseteq s$ and all $d \in D^{\cap}_t$, $t \models_{g[x\mapsto d]} \varphi$

The only novelties in the semantics concern the clauses for the quantifiers. An existential formula $\exists x\varphi(x)$ is supported on a state s in case an element d is established in s for which $\varphi(d)$ is supported. A universal formula $\forall x\varphi(x)$, on the other hand, is supported on a state s if for any possible element d, for any possible enhancement t of s where d is established, t supports $\varphi(d)$.

Thus, a universal assertion $\forall xP(x)$ is supported in s in case it is known not just that P holds of each established individual, but also that it would hold of each possible individual, were it to be established. Analogously, in order for s to support a partition question $\forall x?Px$, it is not sufficient that in s it is known precisely which of the established individuals have the property P and which don't; it must also be established which possible individuals would have the property P and which wouldn't, were they to be established.

Definition 2.6 (Informative content).
The informative content $|\varphi|_g$ of φ w.r.t. g is defined as the union $\bigcup[\varphi]_g$.

Definition 2.7 (Alternatives).
The *alternative set* $[\![\varphi]\!]_g$ of φ w.r.t. g is the set of maximal elements of $[\varphi]_g$. Such maximal elements are called *alternatives* for φ w.r.t. g.

We used to obtain the problematic boundedness formula by a simple existential quantification over an assertion. The following fact shows that this situation cannot arise in the present setting.

Proposition 2.8 (Alternatives for a basic existential).
Consider a sentence $\exists x\chi$, where χ is an assertion that does not contain any constant. Putting $s_d = |d| \cap |\chi|_{[x\mapsto d]}$, we have:

$$[\![\exists x\chi]\!] = \{s_d \mid d \in \mathcal{D}\}$$

Proof. We divide the proof in three parts:

1. for any d, $s_d \models \exists x\chi$;
2. for any state s, if $s \models \exists x\chi$, then $s \subseteq s_d$ for some d;

[3]Notice that this implies in particular that, for $R\vec{t}$ to be supported on s, the elements $[t_1]_g, \dots, [t_n]_g$ must all be established in s.

3. for no $d, d' \in \mathcal{D}$ it is $s_d \subset s_{d'}$.

1. It follows from the definition of assertion that $|\chi|_{[x\mapsto d]} \models_{[x\mapsto d]} \chi$, whence by persistency also $s_d \models_{[x\mapsto d]} \chi$. By definition of $|d|$, $d \in D^{\cap}_{s_d}$. So, the support conditions for the existential $\exists x\chi$ are satisfied on s_d: there is an element $d \in D^{\cap}_{s_d}$ such that $s_d \models_{[x\mapsto d]} \chi$.

2. Suppose $s \models \exists x\chi$. This means that there is a $d \in D^{\cap}_{s}$ such that $s \models_{[x\mapsto d]} \chi$. Now, the fact that $d \in D^{\cap}_{s}$ means that $s \subseteq |d|$, while $s \models_{[x\mapsto d]} \chi$ implies $s \subseteq |\chi|_{[x\mapsto d]}$. Therefore, $s \subseteq |d| \cap |\chi|_{[x\mapsto d]} = s_d$.

3. I still don't know exactly how to prove this formally, but the idea is pretty clear. Since χ does not contain any constant and any free variable other than x, it cannot mention any specific individual. Therefore, given any d and d' in $\mathcal{D}$, the possibilities $|\chi|_{[x\mapsto d]}$ and $|\chi|_{[x\mapsto d']}$ must be perfectly specular. Therefore, they must either coincide, or they must be incomparable.

□

Definition 2.9 (Basic compliant responses)**.**
A basic compliant response to a sentence φ is an assertion χ whose informative content is one of the alternatives for φ.

Proposition 2.10.
For any constant c, the sentence $P(c)$ is a basic compliant response to $\exists x P(x)$.

Proof. It is easy to see that $P(c)$ is an assertion, i.e., it is supported by $|P(c)| = |P(x)|_{[x\mapsto I(c)]}$. Moreover, in order for a state s to support $P(c)$, the element $I(c)$ must be established in s, so it must be $s \subseteq |I(c)|$. In particular, then, $|P(x)|_{[x\mapsto I(c)]} \subseteq |I(c)|$. So we have $|P(c)| = |P(x)|_{[x\mapsto I(c)]} = |I(c)| \cap |P(x)|_{[x\mapsto I(c)]}$. Since the previous proposition ensures that $|I(c)| \cap |P(x)|_{[x\mapsto I(c)]}$ is an alternative for $\exists x P(x)$, we conclude that $P(c)$ is a compliant response to $\exists x P(x)$. □

This illustrates how existentially quantified assertions are generally well-behaved in the system. Of course, this does not yet guarantee that things cannot go wrong for more complex formulas. All I can say is that, for the time being, I was not able to produce a formula that would reproduce the boundedness problem in the present system. One way this could come about is the following: if we can produce an assertion χ whose informative content sets the stage for an (appropriately adjusted) boundedness formula, then the conjunction $\chi \wedge \exists x B(x)$ will not have alternatives. Is this possible? If it is, even if this is theoretically disturbing, can this solution be good enough for any practical purposes? That is, have we pushed troubles away to sufficiently abstruse formulas (such as a conjunction of a long mathematical description with a boundedness formula) that we can be reasonably sure that they will never correspond to anything linguistically relevant? I will leave it at this for the moment.