Yablo's Paradox and the Omitting Types Theorem for Propositional Languages

Thomas Forster

March 7, 2014

This on-line version contains a proof of the extended omitting types theorem which is omitted (Ha!) from the published version

We start by recapitulating Yablo's paradox from [1].

We have infinitely many assertions $\{p_i : \in \mathbb{N}\}$ and each p_i is equivalent to the assertion that all subsequent p_j are false. A contradiction follows.

There is a wealth of literature on this delightful puzzle, and I have been guilty of a minor contribution to it myself. This literature places Yablo's paradox in the *semantical* column of Ramsey's division of the paradoxes into *semantical* versus *logical* paradoxes. However—as I hope to show below—there is merit to be gained by regarding it as a purely logical puzzle.

Yablo's Paradox in Propositional Logic

If we are to treat Yablo's paradox as a purely logical puzzle we should try to capture it entirely within a first-order language with no special predicates. In fact we can even make progress while using nothing more than a *propositional* language; the obvious language $\mathcal{L}$ to use has infinitely many propositional letters $\{p_i : i \in \mathbb{N}\}$. Next we want a propositional theory with axioms

$$p_i \leftrightarrow \bigwedge_{j>i} \neg p_j \tag{1}$$

for each $i \in \mathbb{N}$,

...but of course we cannot do this in a finitary language. However, one thing we can do in a finitary language is capture the left-to-right direction of these biconditionals, and we do that with the simple scheme

$$p_i \to \neg p_j \tag{2}$$

for all $i < j \in \mathbb{N}$.

It can be seen that this is equivalent to the even simpler scheme

$$\neg p_i \vee \neg p_j \tag{3}$$

for all $i \neq j \in \mathbb{N}$.

Let us call this theory Y. Y says that at most one p_i can be true.

It is the right-to-left direction of the biconditionals that gives us trouble ...

$$(\bigwedge_{j>i} \neg p_j) \to p_i \tag{4}$$

for each $i \in \mathbb{N}$.

For each i the right-to-left direction of the ith biconditional (4) asserts that at least one of the formulæ in the set $\Sigma(i)$ is false:

$$\{\neg p_j : j \geq i\} \tag{$\Sigma(i)$}$$

$\Sigma(i)$ is an example of what model theorists call a 0-*type*, a *type* being nothing more than a set of formulæ[1]. The '0' means that the formulæ in the type have no free variables. Our desire that at least one thing in a type should be false is—in the terminology of model theory—a desire to *omit* that type. What we need is a theorem that tells us that a theory can have models that omit a specified type. There is a such a theorem, and it is known as the *Omitting Types Theorem*. We say a theory T in a language $\mathcal{L}$ *locally omits* a type Σ if, whenever $\phi \in \mathcal{L}$ is a formula such that T proves $\phi \to \sigma$ for every $\sigma \in \Sigma$, then $T \vdash \neg\phi$. The omitting types theorem for propositional languages now says:

Theorem 1:

> Let T be a consistent theory in a propositional language $\mathcal{L}$. If T locally omits a type Σ then there is an $\mathcal{L}$-valuation v that satisfies every theorem of T but falsifies at least one σ in Σ.

We say in these circumstances that v **omits** Σ.

However, what we need here is the slightly stronger:

Theorem 2: (Extended Omitting Types Theorem)

> Let T be a consistent theory in a propositional language $\mathcal{L}$. If T locally omits each type Σ in a countable class $\mathfrak{S}$ of types then there is an $\mathcal{L}$-valuation that satisfies every theorem of T but, for each $\Sigma \in \mathfrak{S}$, falsifies at least one σ in Σ.

I will omit a proof of this result, since it is standard in the model-theoretic literature.

[1] A countably infinite set unless otherwise specified.

In asserting the right-to-left directions (4) of the biconditionals we are restricting ourselves to $\mathcal{L}$-valuations that omit all the types $\Sigma(i)$. There are countably many of these types so it would be natural to reach for the extended omitting types theorem, theorem 2. Now if we are to exploit theorem 2 we want our theory Y to locally omit each $\Sigma(i)$. But it doesn't. The formula p_0, in conjunction with the axioms of Y, implies $\neg p_i$ for every $i > 0$ and thereby implies everything in $\Sigma(1)$. If Y were to locally omit $\Sigma(1)$ as we desire then we would have to have $Y \vdash \neg p_0$. But Y clearly does not prove $\neg p_0$. If we were to add $\neg p_0$ as part of a project of adding axioms to Y to obtain a theory that did omit $\Sigma(1)$ we would find by the same token that we would have to add $\neg p_i$ for all other $i \in \mathbb{N}$ as well, and then we end up realising all the $\Sigma(i)$.

Thus Y does not locally omit even one of the $\Sigma(i)$, let alone all of them. So we cannot invoke theorem 2. However, for each i the valuation that makes p_i true and everything else false satisfies Y all right, and it omits all $\Sigma(j)$ for all $j < i$. This illustrates how a theory T can sometimes have a model that omits a type Σ even though T does not locally omit Σ.

Very well: for each i there is an $\mathcal{L}$-valuation that satisfies Y and omits $\Sigma(j)$ for all $j < i$. Can we find a $\mathcal{L}$-valuation that satisfies Y and omits all the $\Sigma(i)$? No! Such a valuation would satisfy all the right-to-left directions of the biconditionals in (1), namely the conditionals in (4) and thereby manifest Yablo's paradox!

Conclusion

Yablo's paradox provides us with an illustration of a setting where there is a theory Y and an infinite family $\{\Sigma(i) : i \in \mathbb{N}\}$ of types where, although Y does not locally omit any of the $\Sigma(i)$, it nevertheless has valuations that omit any finite set of them. Further, it has no valuation that omits them all. That last fact illustrates how the condition in theorem 2—namely that T locally omit every $\Sigma \in \mathfrak{S}$—really is necessary, so the extended omitting types theorem for propositional logic really is best possible.

For T to have a model omitting all the Σ_i it is not sufficient for it to have models omitting any given finite family of them; we really do need the stronger condition that T should locally omit every finite subset of Σ_i.

It illustrates that for T to have a model that omits a type Σ is is sufficient but not necessary for T to locally omit Σ.

This is pædogogically quite instructive!

Appendix: A Proof of the Extended Omitting Types theorem for Propositional Logic

I supply a proof of this fact in this on-line version of the paper because—despite the cheerful observation above that the result is standard in the literature—I cannot actually find a proof anywhere!

A 0-*type* (Hereafter merely 'type': the '0' means that the formulæ in the type have no free variables) in a propositional language $\mathcal{L}$ is a set of formulæ (a countably infinite set unless otherwise specified).

For T an $\mathcal{L}$-theory a T-*valuation* is an $\mathcal{L}$-valuation that satisfies T. A valuation v *realises* a type Σ if $v(\sigma) = \texttt{true}$ for every $\sigma \in \Sigma$. Otherwise v *omits* Σ. We say a theory T *locally omits* a type Σ if, whenever ϕ is a formula such that T proves $\phi \to \sigma$ for every $\sigma \in \Sigma$, then $T \vdash \neg\phi$.

Theorem 1 *The Omitting Types Theorem for Propositional Logic*

Let T be a propositional theory, and $\Sigma \subseteq \mathcal{L}(T)$ a type. If T locally omits Σ then there is a T-valuation omitting Σ

Proof:

By contraposition. Suppose there is no T-valuation omitting Σ. Then every formula in Σ is a theorem of T so there is an expression ϕ (namely '$\top$') such that $T \vdash \phi \to \sigma$ for every $\sigma \in \Sigma$ but $T \not\vdash \neg\phi$. Contraposing, we infer that if $T \vdash \neg\phi$ for every ϕ such that $T \vdash \phi \to \sigma$ for every $\sigma \in \Sigma$ then there is a T-valuation omitting Σ. ■

However, we can prove something stronger.

Theorem 2 *The Extended Omitting Types Theorem for Propositional Logic*

Let T be a propositional theory and, for each $i \in \mathbb{N}$, let $\Sigma_i \subseteq \mathcal{L}(T)$ be a type. If T locally omits every Σ_i then there is a T-valuation omitting all of the Σ_i.

Proof:

We will show that whenever $T \cup \{\neg\phi_1, \ldots \neg\phi_i\}$ is consistent, where $\phi_n \in \Sigma_n$ for each $n \leq i$, then we can find $\phi_{i+1} \in \Sigma_{i+1}$ such that $T \cup \{\neg\phi_1, \ldots \neg\phi_i, \neg\phi_{i+1}\}$ is consistent.

Suppose not, then $T \vdash (\bigwedge_{1 \leq j \leq i} \neg\phi_j) \to \phi_{i+1}$ for every $\phi_{i+1} \in \Sigma_{i+1}$. But, by assumption, T locally omits Σ_{i+1}, so we would have $T \vdash \neg \bigwedge_{1 \leq j \leq i} \neg\phi_j$ contradicting the assumption that $T \cup \{\neg\phi_1, \ldots \neg\phi_i\}$ is consistent.

Now, as long as there is an enumeration of the formulæ in $\mathcal{L}(T)$, we can run an iterative process where at each stage we pick for ϕ_{i+1} the first formula in Σ_{i+1} such that $T \cup \{\neg\phi_1, \ldots \neg\phi_i, \neg\phi_{i+1}\}$ is consistent. This gives us a theory $T \cup \{\neg\phi_i : i \in \mathbb{N}\}$ which is consistent by compactness. Any model of $T \cup \{\neg\phi_i : i \in \mathbb{N}\}$ is a model of T that omits each Σ_i.

■

Obseve that in this above result we do not construct an actual *valuation* of T. What we construct is an extension T' of T with the property that any valuation makes T' true must omit all the types Σ_i. This sounds less useful but it is actually more, for it means that we have actually proved the omitting types theorem for 1-types in LPC as well. The enhancements that follow (using the concept of "locally-∞-omits) are not knwn to generalise to first-order logic.

If $\Sigma = \{\sigma_i : i \in \mathbb{N}\}$ is a [countable] type, then, for each $n \in \mathbb{N}$, the type $\{\sigma_j : j > i\}$ is a **terminal segment** of Σ.

We say a theory T **locally-∞-omits** a type Σ if, whenever ϕ is a formula such that T proves $\phi \to \sigma$ for all but finitely many $\sigma \in \Sigma$, then $T \vdash \neg\phi$. If a valuation v omits every terminal segment of Σ we say v inf-**omits** Σ.

Corollary 1 *Let T be a propositional theory, and $\Sigma \subseteq \mathcal{L}(T)$ a type. If T locally-∞-omits Σ then there is a T-valuation ∞-omitting Σ.*

Proof:

We obtain this from the extended omitting types theorem by thinking of the family of terminal segments of Σ as the family Σ_n in the statement of the extended omitting types theorem. ■

I am skiving out of proving this since there is something stronger that I really am going to prove...

Theorem 3 *Let T be a propositional theory and, for each $i \in \mathbb{N}$, let $\Sigma_i \subseteq \mathcal{L}(T)$ be a type. If T inf-locally-omits every Σ_i then there is a T-valuation -omitting all of the Σ_i.*

Proof: The reader should rehearse the proof of the extended omitting types theorem, and think of it as a construction that, at stage n, picks a representative from Σ_n, negates it and adds it to the theory we are building. The proof of the theorem before us now is obtained from this construction by modifying it to exploit *zigzagging*.

We need a pairing function ($\texttt{pair}(x, y) = (\frac{(i+j)\cdot(i+j+1)}{2} + i) = i$ will do) that admits two unpairing functions, so that $\texttt{fst}(\frac{(x+y)\cdot(x+y+1)}{2} + x) = x$ and $\texttt{snd}(\frac{(x+y)\cdot(x+y+1)}{2} + x) = y$.

We have countably many program variables, $\Delta_n, n \in \mathbb{N}$, and Δ_n is initiated to Σ_n. As in the earlier construction, at each stage we add precisely one formula to a set Φ of formmulæ that we are building (and which is initially empty). At the end of the construction $T \cup \Phi$ will be a theory any valuation for which will ∞-omit every Σ_n.

At stage i we decode i as an ordered pair and look at the first component $\texttt{fst}(i)$, which we will call i'. At this stage we will have already picked $i - 1$ formulæ $\{\phi_0 \ldots \phi_{i-1}\}$ for our set in such a way that $T \cup \{\neg\phi_0 \ldots \neg\phi_{i-1}\}$ is consistent. We want to pick an element ϕ_i of $\Delta_{i'}$ in such a way that $T \cup \{\neg\phi_0 \ldots \neg\phi_i\}$, too, is consistent. If we can find such a ϕ_i we add $\neg\phi_i$ to Φ, and we delete ϕ_i from $\Delta_{i'}$ to obtain the new value of $\Delta_{i'}$.

How can we be sure that such a ϕ_i can always be found? Here we argue much as in the case of the original extended omitting types theorem. Suppose at stage i we could not find a ϕ_i such that $T \cup \{\neg\phi_0 \ldots \neg\phi_i\}$ is consistent, and let i be minimal with this undesirable feature. Then $T \vdash (\bigwedge_{1 \leq j \leq i-1} \neg\phi_j) \to \phi_i$ for every candidate $\phi_i \in \Delta_i$. Now Δ_i is a cofinite subset of Σ_i and, by assumption, T inf-locally omits Σ_i, so we would have $T \vdash \neg \bigwedge_{1 \leq j \leq i} \neg\phi_j$ contradicting the assumption that $T \cup \{\neg\phi_1, \ldots \neg\phi_i\}$ is consistent.

This gives us a theory $T \cup \{\neg\phi_i : i \in \mathbb{N}\}$ which is consistent by compactness. Any model of $T \cup \{\neg\phi_i : i \in \mathbb{N}\}$ is a model of T that ∞-omits each Σ_i.

■

E D I T   B E L O W   H E R E

Now, as long as there is an enumeration of the formulæ in $\mathcal{L}(T)$, we can run an iterative process where at each stage we pick for ϕ_{i+1} the first formula in Σ_{i+1} such that $T \cup \{\neg\phi_1, \ldots \neg\phi_i, \neg\phi_{i+1}\}$ is consistent.

Coda

Since I wrote out this proof some of my enquiries have borne fruit, and I am much endebted to Oren Kolmen for directing me to pp 118–9 of [**?**].

References

[1] Väänänen, Jouko. Models and Games CUP (ISBN-13: 9780521518123)

[2] Steve Yablo, Paradox without self-reference. **Analysis 53.4** (1993) pp 251–52.

Department of Pure Mathematics and Mathematical Statistics
Centre for Mathematical Sciences
Wilberforce Road
Cambridge CB3 0WB
United Kingdom