Nominal techniques and consistency of Quine's NF

Murdoch J. Gabbay

31 May 2016

Introduction

Thanks to the LFCS Seminar series for inviting me here today.
Thank you all for coming.

Bibliography for this talk:

[semooc] "Semantics out of context"
[repdul] "Representation and duality of the untyped lambda-calculus"
[conqnf] "Consistency of Quine's New Foundations"

On naive set theory

Naive set theory has one rule; naive sets comprehension:

- If ϕ is a predicate, then $\{a \mid \phi(a)\}$ (the a such that ϕ) is a set.

This is inconsistent by Russell's famous 1901 paradox:

$$\{a \mid a \notin a\} \in \{a \mid a \notin a\} \quad \Leftrightarrow \quad \{a \mid a \notin a\} \notin \{a \mid a \notin a\}$$

We deduce from this that $\bot = \top$.

Thus 'false is true' follows from the axioms of naive set theory and we can prove anything.

Naive set theory is inconsistent: there are no models of the theory of naive sets, and too many theorems (we can prove anything).

Foundations are fundamental

In the 1930s this was a concern for relatively few people.

Nowadays, thanks in no small part to pioneering work done here at LFCS (Laboratory for Foundations of Computer Science), theorem-provers and programming languages implement these foundations.

Foundations are fundamental

The quality of solutions that were developed to the problem of the inconsistency of naive set theory (see next slide) determines the quality of your working life.

Every time you fire up your computer and write an ML program or an Isabelle proof, as I did during my own PhD, and trust (e.g.) Isabelle not to prove $0 = 1$ and trust ML not to input $1 + 1$ and output 3, then you're living and breathing and trusting in a computer implementation of a mathematical foundation that somebody designed to be powerful enough to be useful but not so powerful as to become inconsistent and prove anything.

Solutions

Solutions proposed:

- Zermelo-Fraenkel set theory (ZF sets).
  Familiar as e.g. "the category of sets", or Isabelle/ZF, and so on.
  'Proved' consistent by the von Neumann cumulative hierarchy model; $\varnothing$, $\mathit{powerset}(\varnothing)$, ...
- Type Theory.
  Familiar as Higher-Order Logic (HOL), ML, and so on.
  'Proved' consistent by taking sets and function-sets; ι, ι^{ι}, $(\iota^{\iota})^{\iota}$, $\iota^{\iota^{\iota}}$, ...
- Quine's New Foundations (NF).
  Restricts comprehension to stratifiable comprehension (more on this in a moment).

NF is pretty

NF admits a universal set: $\{a \mid \top\}$, the set of all sets, is a set.

Unlike ZF and HOL there is no need for an infinite hierarchy of universes/types/classes. In NF $\{a \mid \top\} \in \{a \mid \top\}$, and that's OK!

NF admits the lovely representation of the number n as 'the set of all n-element sets' (due to Frege, 1884).

- Nicer than the standard, brutal, efficient, ZF model: $2 = \{\varnothing, \{\varnothing\}\}$.
- Nicer even than the Church numeral at type α: $2_\alpha = \lambda f{:}\alpha{\to}\alpha.\lambda x{:}\alpha.f(f(x))$.

NF is perfectly safe . . . perhaps!

NF gets a bad rap. It has 'bizzare' properties such as being non-wellfounded and being inconsistent with the Axiom of Choice.

This is unfair: the sets that are non-wellfounded and violate Choice are sets that ZF and HOL do not concern themselves with. Initial segments of the NF universe that look like ZF and HOL exist, and are perfectly definable and usable.

Yes there are monsters in the full NF universe and some people love to go and study them; but these sets won't bother you if you don't actively go looking for them.

Oh . . . but we don't know if NF is consistent. Now that's something to worry about.

NF is pretty

NF's specification is concise; just add 'stratifiable' to naive set comprehension:

- If ϕ is a stratifiable predicate, then $\{a \mid \phi(a)\}$ is a set.

Stratifiability

ϕ is stratifiable when there exists an assignment of an integer level to its variables such that:

- If $a{=}b$ appears in ϕ then $\mathit{level}(a) = \mathit{level}(b)$.
- If $a{\in}b$ appears in ϕ then $\mathit{level}(a){+}1 = \mathit{level}(b)$.

$a \notin a$ is unstratifiable; this blocks the comprehension of Russell's paradox.

$\exists a, b.a \neq b \wedge \forall c.c \in d \Leftrightarrow (c = a \vee c = b)$ is stratifiable, so we can collect all 2-element sets to form

$$2 = \{d \mid \exists a, b.a \neq b \wedge \forall c.(c \in d \Leftrightarrow c = a \vee c = b)\}.$$

Stratifiability

$$Russell = \{a \mid a \notin a\}$$
$$2 = \{d \mid \exists a, b.a \neq b \wedge \forall c.(c \in d \Leftrightarrow (c = a \vee c = b))\}$$

Stratifiability

$$Russell = \cancel{\{a^1 \mid a^0 \notin a^1\}}$$
$$2 = \{d^1 \mid \exists a^0, b^0 . a^0 \neq b^0 \wedge \forall c^0 . (c^0 \in d^1 \Leftrightarrow (c^0 = a^0 \vee c^0 = b^0))\}$$

What is NF?

If we write axiom (**Ext**) and axiom-scheme (**SC**)

$$\begin{array}{ll} (\textbf{Ext}) & \forall a.\forall a.(a = b \Leftrightarrow \forall c.(c \in a \Leftrightarrow c \in b)) \\ (\textbf{SC}) & \exists a.\forall b.(b \in a \Leftrightarrow \Phi) \quad (\Phi \text{ stratifiable}) \end{array}$$

then we can write

$$NF = (\textbf{Ext}) + (\textbf{SC}).$$

That is, NF is extensionality plus stratifiable comprehension.

What a lovely theory!

NF-the-theory vs NF-the-universe

NF is a logical theory, and also a foundational universe.

For the purposes of a consistency proof, NF is a logical theory to prove things about, such as consistency—not a universe to prove things in.

Go catch a very large set of atoms...

If X is a set write $\#X$ for the cardinality of X.

Let $\beth_0 = \#\mathbb{N}$. Write $\beth_\omega$ for the least cardinal larger than $\#\mathit{powerset}^n(\mathbb{N})$ for every $n \in \mathbb{N}$. So:

$$\begin{aligned} \beth_0 &= \#\mathbb{N} \\ \leq \beth_1 &= \#\mathit{powerset}(\mathbb{N}) \\ \leq \beth_2 &= \#\mathit{powerset}(\mathit{powerset}(\mathbb{N})) \leq \cdots \leq \beth_\omega. \end{aligned}$$

Fix a large (size $\beth_\omega$) set of atoms $\mathbb{A}$.

- Write $И a.\Phi(a)$ when Φ holds of all but $\kappa \lneq \beth_\omega$ many atoms. Read this as 'for new a, $\Phi(a)$'.
- Write $Ↄ a.\Phi(a)$ when Φ holds of $\kappa = \beth_\omega$ many atoms. Read this as 'for generously many a, $\Phi(a)$'.

More on this later.

Syntax

Syntactic classes are atoms $a, b, c \in \mathbb{A}$, terms s, and predicates ϕ:

$$\begin{aligned} s, t, u &::= a \in \mathbb{A} \mid \{a \mid \phi\} \\ \phi &::= \phi \wedge \phi \mid \neg\phi \mid \forall a.\phi \mid t \in s \end{aligned}$$

We call $\{a \mid \phi\}$ a comprehension.

Normalise syntax

Now consider the following rewrite on terms and predicates:

$$s \in \{a \mid \phi\} \to \phi[a \mapsto s].$$

Justified by the intuition that s is in the set of a such that ϕ if and only if $\phi[a \mapsto s]$.

Theorem: Stratifiable terms are confluent and strongly normalising under this rule. That is, they rewrite confluently and in finite time to a unique normal form.

Proof sketch: Confluence is routine. Termination follows by rewriting innermost highest level reducts. Use a multiset lexicographic ordering, which is well-founded.

That stratifiability implies existence of normal forms appears to be an original observation of my proof.

Syntax of normal forms

We can easily characterise normal forms:

$$\begin{aligned} s, t, u &::= a \in \mathbb{A} \mid \{a \mid \phi\} \\ \phi &::= \phi \wedge \phi \mid \neg\phi \mid \forall a.\phi \mid t \in a \end{aligned}$$

Note the $t \in a$ on the far right; this is the base case of induction on normalised syntax.

Let a prepoint $p \in \text{Prepoint}$ be a set of assertions of the form $t \in a$. Then we provisionally interpret $t \in a$ by

$$[\![t \in a]\!] = \{p \in \text{Prepoint} \mid (\text{t} \in \text{a}) \in \text{p}\}.$$

Now we want to interpret $\wedge$, $\neg$, $=$, and $\forall$, in the syntax above in such a way as to validate all the axioms of NF.

This will be our model.

Overview of our model

Our model will interpret a predicate as a set of points, where a point is a prepoint plus conditions.

$$\phi \longmapsto [\![\phi]\!] \in \mathit{powerset}(\mathit{powerset}(\{t{\in}a \mid \text{all } t, a\})).$$

So how to interpret logical connectives $\wedge$, $\neg$, $=$, and $\forall$?

Much of this was addressed in [semooc] and [repdul]. I will sketch how it works.

Logic in nominal powersets (propositional part; high-level view)

Conjunction and negation correspond to sets intersection and complement, as usual.

$$[\![\phi \wedge \psi]\!] = [\![\phi]\!] \cap [\![\psi]\!] \qquad [\![\neg\phi]\!] = \mathit{Points} \setminus [\![\phi]\!]$$

(I haven't said which prepoints are points, or proved that any points exist.)

Sets membership becomes substitution, thanks to our rewrite rule:

$$[\![\{b \mid \psi\} \in \{a \mid \phi\}]\!] = [\![\phi[a \mapsto \{b \mid \psi\}]]\!].$$

(This isn't trivial to check.)

Logic in nominal powersets (quantifiers)

What about quantification $[\forall a.\phi]$? Following [semooc,repdul] we write:

$$[\forall a.\phi] = \{p \mid \mathit{И} b.(b\ a)\cdot\phi \in p\}.$$

It turns out that this has many equivalent presentations, including:

$$[\forall a.\phi] = \bigcup\{X' \subseteq [\phi] \mid a\#X'\}.$$

Thus $[\forall a.\phi]$ is the greatest subset of $[\phi]$ for which a is fresh, in the sense of nominal sets.

This characterisation of quantification uses only $\mathit{И}$ and #. It does not depend on substitution!

Logic in nominal powersets (quantifiers)

This

$$[\forall a.\phi] = \bigcup \{X' \subseteq [\phi] \mid a\#X'\}$$

guarantees that:

$$\frac{}{[\forall a.\phi] \subseteq [\phi]} \qquad \frac{[\psi] \subseteq [\phi] \quad (a\#\psi)}{[\psi] \subseteq [\forall a.\phi]}$$

Note we do not use the familiar Tarski semantics that forall = 'for every possible value'. This would read as follows:

$$[\forall a.\phi] = \bigcap_u [\phi[a{\mapsto}u]].$$

That depends on substitution. We can't do that in NF because NF is impredicative and u may be a comprehension $\{a \mid \psi\}$ where ψ is larger than ϕ — taking $\phi[a{\mapsto}u]$ in a definition would be unhealthy for inductive quantities.

Logic in nominal powersets (quantifiers)

The nominal semantics of $\forall$ works generally, just like conjunction and complement.

If $X, Y \subseteq \mathcal{X}$ are subsets of a nominal set $\mathcal{X}$ we can define

$$\forall a.X = \bigcup \{X' \subseteq X \mid a\#X'\}$$

and then

$$\frac{}{\forall a.X \subseteq X} \qquad \frac{Y \subseteq X \quad (a\#Y)}{Y \subseteq \forall a.X}.$$

Thus $\forall a.X$ is the greatest subset of X for which a is fresh.

(This generalises further to nominal lattices; see [semooc].)

Logic in nominal powersets (quantifiers)

Compare with p29 of "Introduction to Categorical Logic" by Awodey and Bauer, where $\varphi \leq B \times A$ and $\vartheta \leq B$:

$$\frac{\vartheta \times A \leq \varphi}{\vartheta \leq \forall_A \varphi} \qquad \text{resembles} \qquad \frac{Y \subseteq X \quad (a\#Y)}{Y \subseteq \forall a.X}$$

They look similar, but there is a crucial difference!

To interpret ϑ and φ we must have objects A and B, and a category including arrows for substitution and so forth. The $\forall_A \varphi$ corresponds to a Tarskian $\bigcap_u \phi[a \mapsto u]$, not to a nominal $\bigcup\{X' \subseteq X \mid a\#X'\}$.

Logic in nominal powersets (quantifiers)

Given an extra consistency condition on prepoints called generous naming of internal sets we obtain a theorem (Theorem 8.15 in the paper):

$$[\![\forall a.\phi]\!] = \bigcap_u [\![\phi[a \mapsto u]]\!].$$

So by the end of my paper, $\forall$ is doing what we expect and quantifying over all terms.

It matters that this is a theorem, not a definition: the ϕ on the right is smaller than the $\forall a.\phi$ on the left in

$$[\![\forall a.\phi]\!] = \bigcup \{X' \subseteq [\![\phi]\!] \mid a\#X'\}.$$

So this is suitable for an inductive definition; the first equality above is not, in NF.

Substitution

An important lemma is that

$$[\![\phi[a \mapsto u]]\!] = [\![\phi]\!][a \mapsto u].$$

This is non-trivial to prove.

Indeed, it is also non-trivial to state. What is $[\![\phi]\!][a \mapsto u]$?

We know that $[a \mapsto u]$ applied to syntax ϕ is.

What is $[a \mapsto u]$ applied to a set of (pre)points like $[\![\phi]\!]$?

Substitution

Suppose $\mathcal{X}$ has a σ-action $x[a \mapsto u]$. Suppose $p \in \mathit{powerset}(\mathcal{X})$ and $X \in \mathit{powerset}(\mathit{powerset}(\mathcal{X}))$ $\mathbf{[}\phi\mathbf{]}$ is one of these X).

Then define:

$$x \in p[u \leftarrow\!\shortmid a] \Leftrightarrow x[a \mapsto u] \in p$$
$$p \in X[a \mapsto u] \Leftrightarrow \not\exists b.(p[u \leftarrow\!\shortmid b] \in (b\ ) \cdot X)$$

Amgis $[u \leftarrow\!\shortmid a]$ is the functional preimage of underlying substitution. The σ-action on X is obtained from the amgis action on p.

Studying the two definitions above is a talk in itself. The bottom line is:

$\mathbf{[}\phi\mathbf{]}[a \mapsto u]$ is obtained by 'lifting' $\phi[a \mapsto u]$ as above.

New and Generous

A filter p generously names x when $\Im a.(a{=}x \in p)$, meaning that $a{=}x \in p$ for $\beth_\omega$ many atoms a.

This guarantees that if $\aleph a.\phi(a) \in p$ then $\phi(a) \in p$ for some a such that $a{=}x \in p$

The mechanics of the proof require the set of atoms to have cardinality $\#\mathbb{A} = \beth_\omega = \bigcup_{i<\omega} \#2^i$.

We unpack this further:

- $\aleph a.\Phi(a)$ holds when $\#\{a \mid \neg\Phi(a)\} < \beth_\omega$.
- $\Im a.\Phi(a)$ holds when $\#\{a \mid \Phi(a)\} = \beth_\omega$.
- $\aleph$ and $\Im$ are dual: $\aleph a.\Phi(a) \Leftrightarrow \neg\Im a.\neg\Phi(a)$.

The technical bits: equality and quantification

The technical rubber hits the mathematical road around Definition 11.19, Proposition 11.30, and Definition 12.37.

We require extensionality, that

$$\forall \phi, s, t, p. \big(p \in [\![s{=}t]\!] \Rightarrow (p \in [\![\phi[a{:=}s]]\!] \Leftrightarrow p \in [\![\phi[a{:=}t]]\!])\big).$$

We enforce this by an inductive construction to build a maximally consistent extensional set of equalities $s{=}t$.

We also require generous naming of internal sets, that for every p and s,

$$\Im a. \forall t. \big(p \in [\![t{\in}a]\!] \Leftrightarrow p \in [\![t{\in}s]\!]\big).$$

We enforce this by another induction generating a maximally consistent set of $(t{\in}a)$s.

The technical bits

Points (well-behaved prepoints) have two faces in my paper.

They have one presentation as $\{t{\in}a \mid (t{\in}a) \in p\}$ and another as $\{s{=}t \mid p \in \mathbf{[}s{=}t\mathbf{]}\}$.

These presentations are equivalent since p 'believes' $t{\in}a$ precisely when p 'believes' $\{b \mid t{\in}a\}{=}\{b \mid \top\}$.

In the latter parts of the paper, we shuttle between the two presentations; one buys extensionality and thus soundness for $=$, the other buys generous naming of internal sets and thus soundness for $\forall$.

About my proof

- Stratifiability gives a normal form for the rewrite $x \in \{a \mid \phi\} \to \phi[a \mapsto x]$.
- $\forall$ modelled using nominal limits.
- Sets extensionality handled by saturating extensionality equalities to a greatest fixedpoint.
- $\forall$-elimination $(\forall E)$ $(\forall a.\phi) \Rightarrow \phi[a \mapsto x]$ modelled by logical dual to И called the 'Generous' quantifier Ↄ.
  Generosity also corresponds to proof-theoretic strength.
- Semantics of predicates as sets of points, where a point is a maximally consistent set of predicates.
- Substitution modelled using $\mathfrak{o}$-algebras.
- Comprehension $=$ atoms-abstractions. So $\mathbf{[}\{a \mid \phi\}\mathbf{]} = [a]\mathbf{[}\phi\mathbf{]}$.
- Atoms extensionally equal to, but not syntactically identical to, comprehensions: $\mathbf{[}a\mathbf{]} = \mathbf{[}\{b \mid b{\in}a\}\mathbf{]}$.
- Two equivalent, but structurally distinct, notions of 'maximally consistent sets': one designed for $=$ and the other for $(\forall E)$.

Conclusions

Obviously, a non-trivial proof. Plenty of moving parts which invite further investigation.

Two threads running through it: those parts specific to consistency of NF, and those having to do with logic and semantics more generally.

Features of the proof include:

- Use of normal forms under rewrite $s{\in}\{a \mid \phi\} \to \phi[a{\mapsto}s]$.
- Analysis of universal quantification $\forall$ in terms of $И$ and # avoids problems with impredicativity and seems to differ significantly from standard Tarski semantics.
- Use of sigma- and amgis-actions on nominal sets. Again, something different.

This slide intentionally not left blank.

(Peek: the sigma-action on syntax)

$$
\begin{array}{lll}
(\sigma\mathtt{and}) & & \mathtt{and}(\mathcal{X})[a{\mapsto}x] = \mathtt{and}(\{X[a{\mapsto}x] \mid X{\in}\mathcal{X}\}) \\
(\sigma\mathtt{neg}) & & \mathtt{neg}(X)[a{\mapsto}x] = \mathtt{neg}(X[a{\mapsto}x]) \\
(\sigma\mathtt{all}) & b\#x \Rightarrow & (\mathtt{all}[b]X)[a{\mapsto}x] = \mathtt{all}[b](X[a{\mapsto}x]) \\
(\sigma\mathtt{eltatm}) & a\#y, x \Rightarrow & \mathtt{elt}(y,a)[a{\mapsto}\mathtt{atm}(n)] = \mathtt{elt}(y[a{\mapsto}\mathtt{atm}(n)], n) \\
(\sigma\mathtt{elt}\mathbf{a}) & & \mathtt{elt}(y,a)[a{\mapsto}[a']X] = X[a'{\mapsto}y[a{\mapsto}[a']X]] \\
(\sigma\mathtt{elt}\mathbf{b}) & & \mathtt{elt}(y,b)[a{\mapsto}x] = \mathtt{elt}(y[a{\mapsto}x], b) \\
(\sigma[]) & c\#x \Rightarrow & ([c]X)[a{\mapsto}x] = [c](X[a{\mapsto}x]) \\
(\sigma\mathbf{a}) & & \mathtt{atm}(a)[a{\mapsto}x] = x \\
(\sigma\mathbf{b}) & & \mathtt{atm}(b)[a{\mapsto}x] = \mathtt{atm}(b)
\end{array}
$$

In $(\sigma\mathtt{elt}\mathbf{a})$, stratifiability ensures definition of substitution is inductive. X may be larger than y, but a' must have lower level than a.

An example: sigma

Nominal algebra (Gabbay & Mathijssen 2006) is like universal algebra but over nominal sets, so enriched with nominal-style names, freshness, and binding.

Let's look at an axiomatisation of substitution $Z[a{\mapsto}X]$:

$$\begin{array}{lrl} a\#Z \Rightarrow & Z[a{\mapsto}X] & = Z \\ & Z[a{\mapsto}a] & = Z \\ a\#Y \Rightarrow & Z[a{\mapsto}X][b{\mapsto}Y] & = Z[b{\mapsto}Y][a{\mapsto}X[b{\mapsto}Y]] \\ b\#Z \Rightarrow & Z[a{\mapsto}X] & = ((b\ a){\cdot}Z)[b{\mapsto}X] \end{array}$$

Call a set $\mathbb{T}$ with an operation σ satisfying the axioms above, a sigma-algebra.

Substitution is an operation of type $\mathbb{T} \times \mathbb{A} \times \mathbb{T} \to \mathbb{T}$ on a nominal set $\mathbb{T}$, satisfying the axioms above.

Just like group multiplication has type $G \times G \to G$ or logical conjunction has type $B \times B \to B$.

Sigma

$$
\begin{array}{lrl}
a\#Z \Rightarrow & Z[a{\mapsto}X] & = Z \\
 & Z[a{\mapsto}a] & = Z \\
a\#Y \Rightarrow & Z[a{\mapsto}X][b{\mapsto}Y] & = Z[b{\mapsto}Y][a{\mapsto}X[b{\mapsto}Y]] \\
b\#Z \Rightarrow & Z[a{\mapsto}X] & = ((b\ a){\cdot}Z)[b{\mapsto}X]
\end{array}
$$

$a\#Z$ is a freshness side-condition. It corresponds to saying 'if a is not free in Z'.

$(b\ a){\cdot}Z$ is a permutation. It corresponds to 'swap b and a in Z'.

Both have natural interpretions over nominal sets.

As lemmas of the concrete syntactic model of syntactic substitution $:=$, the axioms above are often called: garbage-collection, identity, the substitution lemma, and α-renaming.

Cheat-sheet

$$
\begin{aligned}
\text{Stratifiability} &= x \in \{a \mid \phi\} \to \phi[a{\mapsto}x] \text{ terminates} \\
\forall &= И \\
\exists &= Ↄ \\
\text{Extensionality} &= \text{gfp of ``if } x, y \text{ have same elements then add } x = y\text{''} \\
(\forall\mathbf{E}) &= (\#\mathbb{A}{=}\#\text{Sets}) \\
\mathbf{[}\phi\mathbf{]} &= \{p \in \mathit{points} \mid \phi \in p\} \\
\mathbf{[}\phi[a{\mapsto}u]\mathbf{]} &= \{p \mid p[u{\leftarrow}a] \in \mathbf{[}\phi\mathbf{]}\} \\
\text{Sets} &= [\mathbb{A}]\text{Predicates}
\end{aligned}
$$

$\mathbb{A}$ is the set of atoms.