

## Two Negative Readings (of Formality)

### I. Preliminaries

#### A. Review

1. We've started in assessing the formal symbol manipulation construal (FSM)
2. Last time, we distinguished two "readings" of the FSM claim:
  - a. **Positive:** formal systems operate in virtue of the *syntax, form, shape*, or other *potent* property of the ingredient symbols. But we deferred it, to the third critique.
  - b. **Negative:** formal systems work in a way that is *independent of the semantics*
3. Today, push harder on the negative reading: independent of semantics
  - a. Divide it (alas!) into its own two "sub-readings"
  - b. Ask, for each, our standard questions: what does it say, is it true, etc.

#### B. Note

1. From time to time, look at this or that reading, and mention this or that example (such as a theorem prover for a first-order language).
2. Many will have the reaction "*Oh, that's obviously formal!*"
3. May or may not be true; am not willing to say in advance
4. But the issue will be this:
  - a. What does 'formal' mean, such that that system is so obviously formal?
  - b. And can that meaning stand, as the basis of a comprehensive account?
5. My experience is that the surety of intuition that a given system is formal isn't backed up as soundly as one might have thought.
6. So think of this as a kind of detective work—to ferret out from one's deep-seated intuitions what it is that 'formal' could have been meaning, over all these years.

#### C. Motivation (Review)

1. Naturalism
  - a. Why is the negative (independent of semantics) reading of formality so *important*?
  - b. To secure for semantic (*intentional*) systems a degree of naturalistic palatability
  - c. Cf. Haugeland's phrase: the "**make the world safe for semantics**"
  - d. No mysteries: no divine intervention, no behind-the-scenes magic.
  - e. Everything *clear and evident*, laid out in front of us, qua theorists.
2. Status
  - a. This is a legitimate and important role
  - b. If FSM could play this role, it really would deserve its place in intellectual imagination.
  - c. I will argue that it *won't*, that there is no reading that can do this work
  - d. But the *project* shouldn't be forgotten.

## II. Ontological, Conceptual sub-readings

- A. Problem with negative reading is that splits into two *sub-readings* (see figure 1)
- B. First (sub-) reading: **ontological**
- I. Situations in one place don't affect situations in another
    - a. Cf. perihelion of Mercury and the spelling of 'hors d'oeuvre'
    - b. Cf. whether a friend can come for dinner and truth of the continuum hypothesis
  2. Claim:
    - a. This is the reading to which most practitioners are committed
    - b. Hence the one that has figured most prominently in the discussion so far
      - i. Underlying intuition: mixture of *separation* and *non-correlation*
      - ii. I.e., two different readings (i.e., "sub-sub-readings") of "independent"<sup>1</sup>
    - c. Causal "separation":
      - i. We talked about this in Part I, when we were talking about intentionality in general
      - ii. Semantic values—what situations are designated, whether propositions are true, etc. (perhaps even what symbolic structures mean)—are *too far away*, in general, to affect the outcome of the (allegedly formal) inference.
      - iii. How could the state of affairs in a remote region affect what is going on inside the computer (or theorem prover) right here?
      - iv. Examples:
        - α. Far away in *space* (temple at Amritsar, temperature on the far side of the moon)
        - β. Far away in *time* (whether dinosaurs were warm-blooded, what it will be like when the first woman President is elected in the U.S.)
        - χ. Far away in *possibility* (what it would have been like, had there been any bookshelves in the White House during Reagan's tenure)
      - v. Also: not just the semantic value (or designated situation itself), but the facts that *make* it the designated situation
        - α. I.e., the *relation* that the sign or symbol bears to that remote situation
        - β. Unclear what it would be for that relation to have effective force
      - vi. In sum: has to do with the *locality* of symbol manipulation, and the *reach* of semantics (as discussed when we talked about the mind/body problem for machines)
    - d. Non-correlation
      - i. This is an analogous—but not entirely similar—intuition
      - ii. Could *change* the semantics, without *affecting* the syntax
      - iii. E.g., could re-interpret the symbols in a first-order proof of  $p$  and  $p \supset q$  to yield  $q$ .

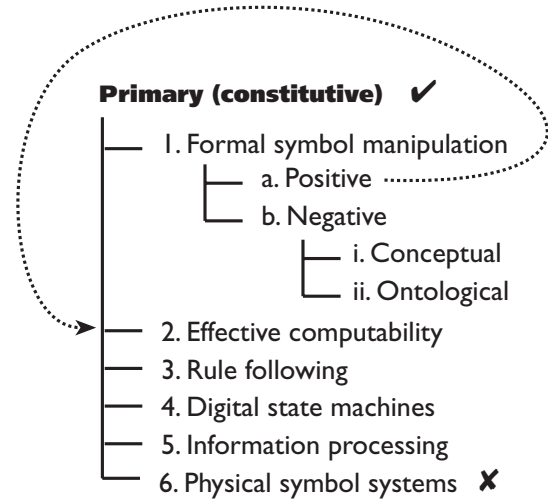


Figure 1 — Two negative readings

<sup>1</sup>We'll come back, in a few weeks, to analyse the notion of *independence* much more thoroughly.

- e. The two sub-sub-readings could part company
  - i. Cf. the numeral–number relation
  - ii. Suppose it were necessary, for some reason (perhaps for unary or binary notations)
  - iii. Might want to
    - α. *Retain causal separation* (because Platonic numbers, sets, functions, etc., cannot cause or influence *anything*, let alone a theorem prover or symbol system); but
    - β. *Discard non-correlation*, because of the presumptively fixed signification relation
  - iv. But then again the two sub-sub-readings aren't wholly independent, either, since non-correlation undoubtedly underlies our intuitions about distance and causality.
  - v. We won't worry about the difference here; run with both, keep distinction in mind.
- 3. In sum
  - a. Characterize this reading as **ontological** or **horizontal**, because of the presumptive “similar level of abstraction” of the two sets of facts.
  - b. Based on an intuition that semantics is “too far away,” in space, time, or possibility
    - i. Situations in one place don't affect situations in another
    - ii. I.e., has to do with *locality* of symbol manipulation, and *reach* of semantics
  - c. Fundamentally: it is a reading based on **disconnection** (an ontological fact)
- C. Second (sub-) reading: **conceptual**
  - 1. Summary:
    - a. Operates at one level higher than ontological story
    - b. *Account* of the syntax and operations can be given independently of the *account* of the semantics
    - c. Because it operates at one level of remove, it is not clear what substantive conditions (if any) it places on the subject matter of computing
    - d. But explanatory role seems familiar enough
  - 2. Underlying intuition
    - a. Can give a precise specification of what happens, how the machine works, etc., without having to describe the semantic interpretation
    - b. I.e., whereas ontological (sub-) reading says that the machine *works* independently of the semantics, the conceptual reading says that those workings can be *described* independently of the semantics.
  - 3. Examples
    - a. (Measure of) length & width of a rectangle: was a good example of ontological independence, but not of vertical
    - b. Better: architecture (of houses): blueprints, independent of description of use.
    - c. Because of meta-level relation, characterise this second form of independence **vertical**
  - 4. Naturalisation
    - a. Fundamentally, this reading is based on a **background naturalisation pressure**.
    - b. A general intuition that somehow or other, formal symbol manipulation either *is* naturalistic already, or helps us see how symbol manipulation more generally could be naturalised—something like that.
      - i. Cf. Haugeland (quoted earlier): “make the world safe for semantics”

- c. I.e., a vague and not-very-well articulated sense—but still a driving intuition—that somehow or other, FSM systems:
    - i. Are themselves naturalistic (or close, or mostly, or ... )
    - ii. May help us see how (serve as a model, etc.) to naturalize others
  - d. This naturalising pressure must be understood to be always in the background
- D. Compare and contrast
1. Quote: second ¶¶ on page AOS II·2·8:
    - a. “Fundamentally, the two kinds of independence, ontological (horizontal) and conceptual (vertical), operate at different levels. The ontological one places the independence squarely in the world: it points at a systematic separation or non-correlation in occurrent phenomena. It is empirically motivated—by the phenomenon of semantic reach, essentially, in conjunction with an enduring sense of no action at a distance. The conceptual one, at one level of remove, locates the independence in the structure of the theory or the types. Its motivations are theoretic, having to do with autonomy of analyses. To assume that they are the same is simply a confusion; to assume that they align (a more interesting mistake) is to assume that the joints of the world and the joints of the theory line up—a remarkable, and ultimately untenable, assumption of word–world isomorphism.”
  2. Actually, a third notion
    - a. **Explanatory independence:** stemming from naturalism: the form that reductionist theories have, on the right hand side of their equations
    - b. Characterisation *free of undischarged semantical predicates*
    - c. Cf. “hierarchy of nature”
    - d. So there is a “three-way thicket to untangle”
  3. Summary
    - a. *Disconnection* and *naturalisation* both lead to independence claims, but they are in *tension*.
- E. State of the art
1. Ontological reading: won’t satisfy the explanatory requirement
    - a. Just claims there are two phenomena going on (one syntactic or proof-theoretic or “computational”, the other semantic or intentional or referential)
    - b. Quote: from ¶¶ 2 on page II·2·10:
      - i. “Ontological formality characterises computation in semantical terms, albeit negatively, claiming that from an operational point of view the phenomenon works without it. What would answer the reductionist’s dream is exactly the opposite: a non-semantical (naturalistic) characterisation that shows how actual (not banished) semantics genuinely comes about. It is all reminiscent of a simple use-mention confusion. The desired independence is at the level of explanation; what we are offered is at the level of the phenomenon. And the substitution does not fly.”
  2. Neither will conceptual (on its own)
    - a. Just because the semantic story is *separate*, doesn’t mean that it won’t have to be *given*
    - b. Cf. one-two punch (semantic implementation):
      - i. Give an account of the syntactic (in naturalistically palatable terms)

- ii. Then, give an account of the semantic, grounded in the syntactic
  - iii. I.e., define syntax in terms of mechanism (or physics), semantics in terms of syntax.
  - c. But success of that scheme *depends on the falsehood of the horizontal claim!*
    - i. Semantic implementation would work *only if separation/non-correlation were false*
    - ii. One of the great ironies!
    - iii. Substance of the horizontal claim: *syntax and semantics are different things, at the same level of analysis*
    - iv. Hence its label of 'horizontal'
  - d. What would satisfy the naturalistic program is the opposite:
    - i. Syntax and semantics would need to be the *same thing, at different levels of analysis*
    - ii. Hence its label of 'vertical'
    - iii. Quote: bottom ¶ on page II·2·11:
    - iv. "To the extent that the horizontal reading is true, therefore (and it does rest on a genuine insight, even though I will argue that in idealised form it fails as a constitutive characterisation of computation), it follows that the vertical reading—to the extent that *it is true*—fails to satisfy the naturalistic requirement of providing a palatable account of the constitutive conditions for being intentional. The very real grain of truth underlying the ontological (horizontal) reading of formality, that is, blocks the reductive success of the conceptual (vertical) program, with the net result that *all the constitutive semantical questions are left open.*"
- F. Summary
- 1. In the end—when all is said and done—will be impossible to credit the idea that the notion of formal symbol manipulation can simultaneously discharge the intentional mysteries and mesh with other (allegedly more acceptable) forms of understanding of the world.
  - 2. That illusion will have been shown to rest a welter of conflated distinctions and scope ambiguities, having to do, among other things, with:
    - a. How the general case is derived from the particular;
    - b. What aspects of full intentional situation are dealt with by various different explanations;
    - c. What senses of independence are in play at what point, and how they relate;
    - d. What the proper relation is between theoretic, epistemic, and ontological conditions;
    - e. How notions of effectiveness, physicality, and causation tie together (difficulties endemic to the cluster of potency predicates); and
    - f. How abstraction and implementation boundaries relate to issues of theoretic reduction

### III. Catechism

- A. So far, this is all pretty abstract.
- B. To make it vivid, consider a familiar case
  - 1. Language L, variables **x**, **y**, **z**, etc., forms of composition, inference rules, etc.
  - 2. Two stage story:
    - a. Syntax and operations
    - b. Semantics
  - 3. Show the "syntax mirrors semantics"

- a. Not how the semantics *arises*
  - b. Or: *what makes the intended interpretation be the intended interpretation*, etc.
  - c. Rather: given the semantics, how the coordination (between syntax and semantics) goes
4. In sum:
- a. Semantics is *assumed*
  - b. What the soundness, completeness results etc. show is *that the operations honour them*
  - c. So the part that was mysterious (semantics) is not explained.
5. Cf. voodoo doll (story from chapter)
- C. Possible objections
1. Semantic absorption
  2. Inappropriate generalization
    - a. One of the sorts of worry we will have is that some condition is *built in* to the kinds of languages we study, so as to ensure that the way they are described works.
    - b. for example, suppose that **P** means “**P** has not yet been explicitly denied,” then ‘**P**  $\vee$  **Q**’ could be true without warranting an inference from ‘**P**  $\vee$  **Q**’ and ‘ $\neg$ **P**’ to ‘**Q**’.
    - c. But that raises all kinds of questions;
      - i. What is that condition?
      - ii. Why is it built in?
      - iii. What consequences does it have?
      - iv. Is it true?
  3. Construction
    - a. We usually *specify* the semantics; don’t *describe* or *discover* them
    - b. Recipe for building, not theory for analysing?
- D. Suggestion
1. Catechism has the form that *both stories are simultaneously true*
  2. This may be the most telling fact of all.
  3. Is it possible that what it is to be a formal system is to be such that *both ontological* (horizontal) *and conceptual* (vertical) independence *are true at the same time*?
  4. A kind of “alignment” suggestion, for the two (sub-) readings?
    - a. Quote: final ¶¶ on page II·2·27:
    - b. “This suggestion is not completely bizarre. When two issues are genuinely ontologically distinct—such as the mathematical question of whether  $\gamma$  is irrational, and the social question of whether the theatre critic you ran into last night should be invited to the cast party—they naturally submit to independent explanations. On the other hand, it seems odd to argue that syntax and semantics are *that* independent. In fact the trick, regarding syntax and semantics, is how to keep them far enough apart to retain the alleged advantages of formality (specifically, of satisfying some form of naturalistic yearning, and of avoiding the most problematic or mysterious aspects of semantics), and yet close enough together to avoid a “magic coincidence” or dualist conception of parallel universes. That is what is hard, and that is what makes the alignment suggestion seem odd. For it seems *too* independent.”

#### IV. Range, Level, and Aspect

- A. Try to distill what is going on
- B. Five criteria that an explanation or theory should meet:
  - 1. Three are empirical
    - a. **Range**
      - i. It should apply to all possible types (and instances) of the subject matter, not just one specific architecture or kind;
    - b. **Aspect**
      - i. It should explain all constitutive or essential properties or attributes, not merely a restricted subset or projection; and
    - c. **Level**
      - i. It should treat these properties or aspects at the level at which the salient generalisations or regularities hold, not (or at least not only) at an underlying level of implementation.
  - 2. Two are meta-level, on the form of a theory
    - a. **M-Explanation**
      - i. It should offer a substantial and illuminating explanation of the phenomenon in question, not just point towards or name it.
    - b. **M-Naturalisation**
      - i. It should help us understand how the phenomena in question are “natural” (not mysterious, divine, or otherwise metaphysically spooky)
      - ii. One (but only one) way to do this is by *reduction*: to formulate an account of the phenomenon in question (“on the right hand side of the equations,” as it were) in language free of undischarged semantical predicates.
- C. Explanation
  - 1. On **m-naturalisation**
    - a. Note that eliminativism doesn’t meet **m-naturalisation**
      - i. **M-Naturalisation** wants you to eliminate semantical predicates from the *right*
      - ii. Eliminativism eliminates them from the *left*
    - b. Point of **m-naturalisation** is to *show how semantics, interpretation, etc. are part of nature*, not to show that there *isn’t any such thing as semantics, interpretation, etc.*
- D. Eight important examples (see figure 2, at the top of the next page)
  - 1. Full account of the possible syntax and operations of all possible computer systems (e.g., to peek ahead to the second construal)
    - a. Meets **range**
    - b. Fails **aspect**
    - c. May be okay on **level**
  - 2. Complete analysis of the human condition, including an account of our values, our interpretations, our semantics, our consciousness
    - a. Meets **aspect**
    - b. Fails **range**

	range					
	aspect					
		level				
			m-explanation			
				m-naturalisation		
1	✓	✗	(✓)	(✓)	(✓)	Operations & behaviour (mechanism) of all possible computers (cf. ec)
2	✗	✓	(✓)	✓	?	Complete analysis of human condition (emotions, consciousness, values)
3	✗	✗	✗	✓	✓	Full theory of the neuroanatomy /physiology of human brain (cf. ucsc)
4	...	✗?	...	...	(✓)	Semantic eliminativist (someone who denies there is semantics or meaning)
5	✗	✗	✗	(✗)	(✗)	Builders of the first genuinely successful AI robot (cog?) _ worst case
6	✗	✓	✓	✗	✗	Traditional proof and model theory (trad. catechism for first-order logic)
7	?	(✓)	(✓)	✗	✗	Ontological (horizontal) reading of fsm
8	?	(✓)	(✓)	✗	✗	Conceptual (vertical) reading of fsm (m·nat ok semantic implementation)

3. Full theory of the neuroanatomy and neurophysiology of the human brain (cf. UCSD)
  - a. Could fail to meet *all three!*
4. Semantic eliminativist
  - a. May meet **m-naturalisation**, but not by *reduction*
  - b. **M-Reduction** wants you to eliminate semantical predicates from the *right*
  - c. Eliminativism eliminates them from the *left*
  - d. Point of **m-reduction** was to combine **level** with **m-naturalisation**
    - i. I.e., show how semantics, interpretation, etc. are part of nature
  - e. Elimination does the opposite:
    - i. Not *show how* semantics, interpretation, etc. are part of nature
    - ii. But *deny that* semantics, interpretation, etc. are part of nature
5. Construction (success in building some certain machine)
  - a. Could fail on **range**, because only a single exemplar
  - b. Could fail on **aspect**
    - i. Cf. money, stop signs
    - ii. Don't need to *build* constitutive social (relational) facts, even if they are necessary.
  - c. Could (likely will) fail on **level**, because you build things at a lower level of description than that at which they are instances of the type of interest.
6. Catechism (for first-order logic):
  - a. Fail on **range** (because so limited—cf. diagrams, non-monotonic and unsound, etc.)
  - b. Certainly do better on **aspect**
    - i. Succeed, even, if inference and interpretation exhaust the salient qualities
    - ii. This is why I am so interested in logic, of course: it recognizes the importance of the semantical (a condition that **aspect** ferrets out)
  - c. Succeed on **level**



- d. Fail on **m-explanation**
    - i. This is the problem with semantic absorption
    - ii. (Well recognized in philosophy of mind)
  - e. Perforce, fail on **m-naturalisation**
    - i. Succeeds in not having any undischarged semantical predicates on right hand side.
    - ii. But does so *by not having any right hand sides at all*
    - iii. Somehow that doesn't count.
7. Ontological (horizontal) sub-reading of negative reading of FSM construal
- a. Fails **m-explanation**—because no explanation is given
  - b. Fails **m-naturalisation**
    - i. Because identifying filter (“without regard to semantics”) is intentionally defined!
    - ii. Extremely important: quote (end of page II·2·33):
    - iii. “With respect to meeting the naturalistic demand for an ‘independent of semantics’ account, *it makes not a whit of difference whether computers do or do not make use of their assumed-to-be-genuine semantical relations*. Banishing semantics to the wings—especially banishing it *qua* semantics—doesn't even *begin* to pay the naturalist's debt.”
  - c. **Aspect**—could succeed, but fails to the extent that it puts emphasis on only syntax
  - d. **Range** — hard to say (that is what we want to ask next time)
  - e. **Level** — maybe okay.
8. Conceptual (vertical) sub-reading of negative reading of FSM construal
- a. **Range** — hard to say (that is what we want to ask next time)
  - b. **Aspect** — pretty good (again, recognizes the centrality of the semantic)
  - c. **Level** — seems okay
  - d. **M-Explanation**
    - i. Doesn't really do anything
    - ii. Neither solved nor foreclosed (because it is a meta-level condition)
  - e. **M-Naturalisation**
    - i. This is most interesting (because motivation), but the result is hard to characterize.
    - ii. Depends, for one thing, on whether the “semantical” predicates that the proposed syntactic theory (**t-syntax**, in the book) can be described independently of are the only naturalistically challenging ones.
    - iii. I.e., a question of whether the “syntax inheres in the physics.” Cf. Searle
    - iv. Also depends on whether the semantics inheres in the syntax
      - α. If so, then **t-semantics** would satisfy **m-naturalisation**
      - β. But semantic implementation is *false*
      - γ. As we said before, that is guaranteed by the (considerable) truth underlying the horizontal reading.
      - δ. So nothing in this reading puts any condition on **t-semantics** satisfying **m-naturalisation**
    - ε. It still could, but that is a separate fact, that would have to be secured through other means.

## E. Summary

1. Time to let go
2. **M-Naturalisation** won't survive (at least in my hands)
3. Immediate question (take it up next week) is whether either of these readings (horizontal or vertical, ontological or conceptual), is actually (empirically) *true*.

— end of file —