

Independence

I. Preliminaries

A. This is the final week of our analysis of the FSM construal

B. Last time, we completed a first pass through the empirical critique

1. Started out with a detached, isolated, and aligned picture
    a. Pure **internal world** of abstract digital symbols, manipulated on their own, with minimal dependence (causal, semantic, explanatory) on the outside world
    b. Potentially messy, continuous, concrete **external world**, to which the symbols refer, but not in a way that affects their treatment
    c. *Boundary* between the two mediated by **transducers**—all-purpose devices for crossing allegedly aligned boundaries
2. Ended up with quite a different picture
    a. Important: on the positive reconstruction of this construal, *we still envisage a realm of symbols, with syntactic and semantic aspects*
    b. But (like Joshua) we blew our horns: and the boundary between the machine and the world came tumbling down.
    c. Unearthed a spate of important and interesting distinctions
        i. Between in the inside of a computer and the outside
        ii. Between symbols and their referents
        iii. Between continuous and discrete phenomena
        iv. Between and among various degrees of abstraction
        v. Between one medium and another
        vi. Between (static) programs and the (dynamic) processes they engender
    d. Overall a *participatory* world: **computers are involved in their subject matters**.
    e. We will keep this conception of an inexorably participatory semantic realm to the end.

C. Plan

1. This "middle region" of participation is a non-trivial result. So are we done? No.
2. As promised, we need to go back to pick up the thread we left lying on the ground
3. For a schematic summary of the state of the investigation, see figure 1, on the next page.
4. When we came into the FSM construal and critique, we made the first major distinction: between **positive** and **negative** readings of the underlying intuition:
    a. *Positive:* works in virtue of syntactic, formal, effective, causal properties
        i. Called this the issue of **potency**—of *how machines work*.
        ii. Didn't have much to say about it, how. We deferred to the second construal (EC)
    b. *Negative:* overarching banner is of **syntax independent of semantics**.

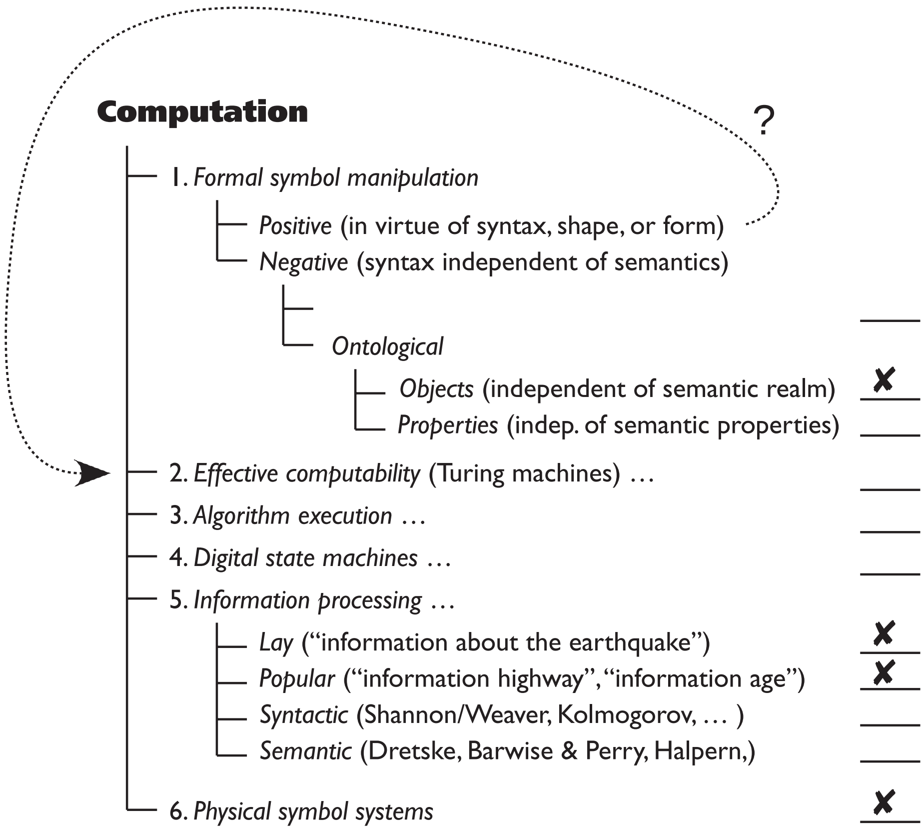

Figure 1 — State of the Investigation

5. So far, we have pursued the negative reading, which then split into two *sub-readings*:
   a. **Conceptual**: an *account* of how the machine works without reference to an account of what it refers to or how it is interpreted
   b. **Ontological**: the machine actually *works* independent of the semantic properties
6. In turn, we split the ontological reading, in turn, into two more *sub-sub-readings*:
   a. Phrased at the level of **objects***:* that computing works *independent of the semantic realm*
      i. I.e., independent of the task domain
      ii. This thesis is false (this is what we've argued).
      iii. On the contrary, far from operating independent of their semantic realms, computers are involved in those realms (i.e., in their subject matters).
   b. Phrased at the level of **properties**. Argument goes as follows:
      i. Even if semantic values (referents) are causally (effectively, potently) implicated in the workings of the machines, and
      ii. *Even if relations between symbols and their referents* play a constitutive role in the workings of the mechanisms (that is, the mechanism wouldn't "be" the mechanism that it is—say, wouldn't have been built),
      iii. Still, they do so in virtue of *being causal relations*, not *being semantic relations*;
      iv. So computers work independently (w/o regard to) semantic properties, after all.

7. We need to examine this property version next
8. I will argue that there is:
   a. (✔) Something right about it; and
   b. (✘) Something wrong about it.
9. There are three complexities, en route
   a. Substantive
      i. NB: although I will deny that the property reading is true (i.e., will deny that it works as a constitutive account of what it is to be computational), *I will not argue the converse* (i.e., will not argue that the way in which computers work is *dependent* on the exemplification of semantic properties)
      ii. That is:
         α. I will deny that syntax (effectiveness) and semantics are 90° apart
         β. But I will not argue that they are 0° apart, either — i.e., that they are the same.
      iii. Rather, I will formulate and defend a thesis of *partial interdependence*.
   b. Methodological note #1:
      i. This is very tricky stuff
      ii. Any attempt to make sense of the issue inevitably trips over all sorts of tough questions in metaphysics, ontology, and the philosophy of science. (It's probably because of these difficulties that the situation has never before been adequately clarified.)
      iii. Reason: because we are going worry about four complex, intertwined topics:
         α. Property identity
         β. Varieties of modal dependence (logical, conceptual, empirical …)
         χ. Issues of reduction and supervenience
         δ. The notions of dependence and independence themselves
      iv. The first of these (property identity) is hard exactly because *objects* are the paradigmatic type of thing over which identity is defined!
      v. So dealing with the property formulation of the antisemantical argument involves a kind of "abstraction ascent": the objectification of properties. This is a dicey thing to do with any certainty.
   c. Methodological note #2:
      i. We will be trying to make these intricate manoeuvres *sans a theory of either syntax or semantics*. In a way, this may seem crazy; one might expect that we shouldn't wrestle with such questions until we have a candidate theory to contend with.
      ii. Ironically, however, I believe it is essential to do now—in order to develop our intuitions about what syntax and semantics are (actually) like
      iii. So will do the best we can.
10. In spite of these difficulties and methodological challenges, we will be able to do two things:
   a. Develop a reconstructed formulation of the essential FSM insight (which we haven't uncovered, yet) that will serve us in a positive account;
   b. Along the way, learn some more things about the nature of computing and semantics.

II. Summary

A. Argument is tough enough that it will help to summarize it first, before making it.
B. First, summarize the challenge (this is extracted from book chapter: AOS·II·5 – §1)

Some will say that this entire analysis is flawed. The problem, they claim, stems from an insufficient distinction between objects and properties they exemplify. Antisemantical formality, they argue [i.e., the "independent of semantics" version that we have called the "negative" reading of formality], even on the ontological (horizontal) reading, was never meant to imply that *referents* are forbidden from playing a causal role in engendering behaviour. Admittedly, in the normal examples—first order logic, Platonic arithmetic, NASA systems calculating planetary trajectories, and the like—symbolic and referential realms are effectively (and explanatorily) separated, and referents therefore never get into the causal act. But that merely simplifies the examples, making the point more accessible to the imagination. The substance of the ontological thesis, or so at least the critics will claim, is not about objects *per se*. Rather, they will argue, what the formality condition mandates is that the behaviour of a computer must (positively) proceed in virtue of the exemplification of syntactic or other potent properties, and (negatively) *not in virtue of the exemplification of semantic properties*.

Thus suppose, as in the case of quotation, counting, electronic mail, internal reference, embedded interaction with the environment, and the like, that some sign or symbol α not only refers to β, but that β also plays a causal role in the treatment of α. From these facts alone, the critic would argue, there is no reason to indict the antisemantical thesis [i.e., negative reading of formality]. For β plays that causal role because of its syntactic (effective, causal) connection with α. The fact that it *also* happens to be α's referent is causally irrelevant. One can show that it is irrelevant with a simple thought experiment: simply change the semantics, and reinterpret α as designating γ. Would the behaviour change? Absolutely not; β would go on playing exactly the same role as before. This proves, the critic claims, that β does not play a causal role in the treatment of α in virtue of being α's referent, but rather in virtue of being, in some way or other, *physically connected* to α

In sum, the critic concludes, every one of the examples of internal reference is misleading. In each case one object (β) is characterised by a **coöccurrence** of properties: first, of being the referent of some other object (α); and second, of figuring in the effective treatment of that same other object. The length of '"abc"', for example, is not only a semantic property (being the referent of 'length("abc")'), but also a syntactic property. The thought experiment shows that the two properties are different, and (tautologically) that it is the latter that affects behaviour. So the examples actually support, rather than contradict, the original claim that computational processing is semantically blind.

In making this critique, the critic can even admit that coöccurrence is theoretically striking, and perhaps unexpectedly common. They just will not see it as undermining the antisemantical claim. It is perfectly ordinary, the critic will claim, for different properties to enter into different regularities, including properties that *overlap in extension*. This is one of the reasons to push for a property-level formulation. For essential characterisations must always be phrased in terms of underlying, constitutive properties—must be phrased, that is, *intensionally*. And once one recognises the fine-grainedness of the constitutive regularities, one will see, according to the critic, that causal properties enter into behavioural consequence; semantic ones, into intentional interpretation. The two, in this way, remain distinct, and formality triumphs after all

This is an important criticism. It fails, ultimately—in the sense that it does not finally rescue the antisemantical thesis. Nor does it do anything to lessen the theoretical impact of the thick, participatory interplay that has been uncovered between symbolic and semantic domains. All the morals we have been elicited so far—about the non-alignment of physical and semantic boundaries, about ambiguities surrounding the cluster of potency predicates, about the participatory (rather than deductive or experiential) nature of computation, about the ambiguous status of the notion of a transducer, about the "middle region" locus of interesting computational structures, and about the question of how computers are to be individuated—all these results would stand even if the coöccurrence critic were right. Indeed, one reason for having analysed an informal, object-level version of the ontological thesis, rather than having moved directly to the more complicated version in terms of properties, was in order to elicit this much of a positive, participatory picture of what computation is like.

Nevertheless, there is merit to the challenge. It is right, for example, to insist that the discourse be conducted intensionally, in terms of properties, rather than purely extensionally, in terms of objects. For it is properties, after all, or at least what we get at by registering the world in terms of properties, and not the objects that exemplify them, that figure in nature's laws and regularities. Moreover, responding to the fine-grained criticism will give us our first taste of what semantic relations are actually like—an issue on which the discussion has so far remained inauspiciously silent. In addition (something we haven't seen before), by shifting attention from objects to properties, the criticism will push the ontological (horizontal) reading of antisemantics right up against the conceptual (vertical) reading, which makes addressing it a profitable exercise before turning attention to the latter topic.

C. So that's the challenge: what we can call the **challenge of the coöccurrence critic**

D. In a moment I will summarize my reply. But first, we need to identify three kinds of people, who may surreptitiously try to hitchhike a ride with the coöccurrence critic, but must not be allowed to distract the analysis (we will have to be on the lookout for them):

1. **Eliminativists**
    a. Some people will say that *of course* a computer cannot depend on the semantics, because there aren't any semantics to depend on.
    b. As said before: such people cannot be advocates of the FSM thesis at all.
    c. To say that would be "independent of semantics" in the sense that physicists wanted physics to be *independent of the luminiferous ether*.
    d. What we want, rather, is "independent of semantics" in the sense that prohibitionists extolled the virtues of lives *independent of alcohol*.
    e. So eliminativists have no bite here. (They don't have stage passes to this concert.)
2. **Reductionists**
    a. Others might say that a computer cannot depend on the semantics, because "all semantics is is α," for some α—and then show that α is not semantical (because it is just causal, say).
    b. This again confuses the kind of independence we are considering
    c. Specifically, what they are concern about is explanatory independence (of the sort motivated by the naturalistic tendencies we saw earlier)
    d. Compare:
        i. "My house wasn't burned down by the *fire*, it was just burned down by the oxidation of plant material"; or

ii. "It wasn't the *heat* that burned you; you were burned by the root mean square velocity of the air molecules.

e. In both cases, that claim is wrong:
   i. Suppose that fire just *is* oxidation, or that heat just *is* RMS molecular velocity.
   ii. Then the house *did* burn down because of the fire, if it burned down in virtue of the oxidation of plant material, because that is just what fire *is*.
   iii. And you *were* burned by the heat, if you were burned by RMS molecular velocity—because, again, RMS molecular velocity is what heat *is*.
f. Moral: in the face of constitutive participation in the task domain, you can't secure the truth of the FSM thesis through reductive or eliminativist manoeuvres.
g. If you accepted a pass to this party, you must have agreed that there *is* semantics, for computation to proceed independently of—whether that semantics is intrinsic (original) or attributed (derived).

3. Quick change artists

a. Some people (e.g. Searle, coöccurrence critic) may try to argue by saying:
   i. "I can re-interpret the wall in front of me as an implementation of WordStar"; or
   ii. "You claim that this program is proving theorems of arithmetic, but suppose I reinterpret it as deciding the fate of six hostages I am keeping in the basement."
   iii. In neither case, the critic claims, will the machine do anything different.
   iv. So that proves that semantics is irrelevant to behavior
b. Reply, part #1: individuation
   i. Consider: *the behavior would not be any different*, under this change of interpretation
   ii. The change *is* relevant to behavior, *if behavior is individuated semantically*
   iii. For example: the behavior of the system used to be *proving theorems of arithmetic*, whereas now (according to you) it is *deciding the fate of hostages.*
   iv. Viz.: different behaviour! So you *have* changed it, after all.[1]
   v. Clearly, this was not what the critic (or Searle) meant
   vi. Rather: what must be intended is something like "same behavior *under syntactic or effective description*," something like that—i.e., same behavior *causally individuated.*
   vii. So some notion of syntax and/or effectiveness built into the challenge.
c. Reply, part #2: arguing to the premise
   i. Who says you can just reinterpret the system in that way?
   ii. Maybe you would be wrong!
   iii. By analogy: suppose, in a discussion of whether cats use maps of their territories as a way of getting around. In arguing *no*, someone says: "Look, I can interpret the cat's patterns of neuronal activity as a representation of Bach's *St. Matthew's Passion*, not as a map. Under that reinterpretation, there is no longer a map in the cat's head. And the cat continues to catch mice and find its way home fine. So cats cannot be using spatial representations."

---

[1]We will see lots of problems of this kind—what happens to individuation when you allow semantical properties to play a role—when we consider Turing machines, in part III of the course.

iv. That's a crazy argument, and flawed for the same reason it is flawed in our case.
v. You can't just *assume* that the semantics can be changed at will, by a (causally) external observer (the way Searle does)
vi. To do that (given the above remark about behaviour being individuated causally) is simply to *assume* an independence of the two.
vii. But that was what we were supposed to be showing!
viii. In general, people who think that computers are such that you can change the semantics at will are people who are committed (on at least some reading) to the FSM view. To be committed to it *in advance* is just ideology.
ix. What we were supposed to be doing was to see whether the FSM view is true.

d. In general, I will dub people who assume that you can change the semantics or interpretation of a system at will—thinking that they can use this as part of argument for the formality of computing—**quick change artists**.

4. In sum: neither **elimination**, nor **reduction**, nor **quick changes**, are allowed in court.
5. Question is whether the independence thesis can survive, without their aid.
6. Answer: *no*. This is what I now want to argue.

E. Summarize the reply

1. We are using, as a basic assumption (this is why we pursued the "object" analysis first)

> ◆ **Involvement:** Computers are *involved* in their subject matters.

2. Because of INVOLVEMENT, the relation between syntactic and semantic properties is *partially constitutive*, not *purely coincidental*
   a. True on various readings of what it is that establishes the semantics:
      i. **Realist** (those who claim there is a fact of the matter as to what the semantics are): facts about the system, as embedded in an environment, help to determine what the semantical values are—be those facts historical, evolutionary, social, physical (whatever you think is crucial).
      ii. **Interpretationalist** (those who think semantic values are attributed from the outside): there are constraints on what we can take systems to mean or represent—cf. the "principle of charity", the end of Haugeland's introduction to *Mind Design*, etc.
      iii. Some admixture of the two—or if, like me, you don't think the distinction is a genuine one (see recent Dennett).
   b. Enough to defeat the independence claim (FSM failure #1: *falsehood*)
   c. But not enough to extract the insight out from the original FSM thesis
3. More specifically: semantics is in part *constituted* by syntax and behaviour.
   a. Not enough to preserve FSM (FSM failure #2: *blunt instrument*)
   b. But yes, enough to extract the fundamental FSM insight
4. To get at that insight, need to tease apart the independence claim.
   a. Specifically, distinguish two claims:

◆ **Independence:** Change the semantics without changing syntax
◆ **Impotence:** Semantical properties are not causally effective

b. IMPOTENCE is the one for which we have basic ontological motivation (cf. what is right about disconnection).
c. But INDEPENDENCE is what embodies (the property reading of the ontological reading of) the independence claim.
d. Might think that IMPOTENCE would entail INDEPENDENCE. **But it is not so!**

5. Finally we are getting results:
   a. Result #1: INDEPENDENCE will *not* survive, because of essential relation between syntax (effectiveness) and semantic
   b. Result #2: IMPOTENCE *will* survive (at least we will make room for it, allow it to survive for now; and sure enough, once we are done, we will see that it survives to the end).

F. Strategy
   1. Could get at these results by going straight after syntax and semantics: what they are like "on their own," as it were.
   2. Problem: we don't have theories of syntax and semantics strong enough to do that!
   3. Instead: examine notions of dependence and independence. Because the issues have fundamentally to do with how syntax and semantics *relate.*
      a. Independence: centerpiece of FSM claim!
      b. But there are different senses: extensional, intensional
      c. Same for *d*ependence: different flavours—causal, constitutive, etc.
      d. Only when these different notions are sorted out can real structure of the FSM situation shine through.
   4. So: analyse dependence and independence, as organisational strategy

G. So turn to independence. Like so many other things, *it too comes in multiple forms.*

III. Extensional independence

A. Intro
   1. **Extensional independence:** a simple matter of more or less correlated variation in the world
   2. Cf. figure 2 (to the right)
   3. Attempt a definition:
      a. A set of two or more entities {$\mathbf{x}_1$, $\mathbf{x}_2$, … $\mathbf{x}_k$} are **extensionally independent** just in case, with respect to (i) ranges of potential variation or sets of possible choices of elements {σ_1, σ_2, … σ_k} of each, and (ii) a back-

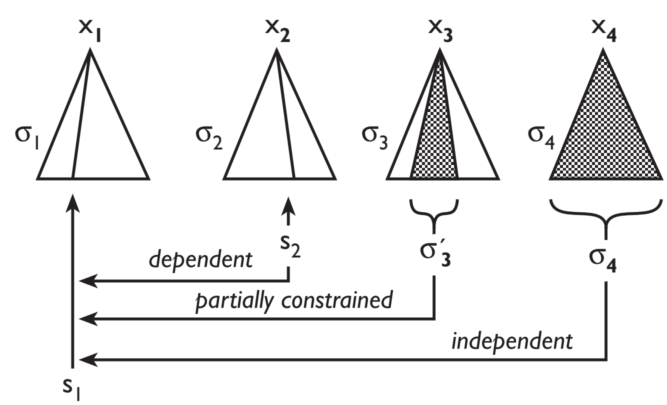

Figure 2 — Dependence, independence, and constraint

ground condition **c** that must be honoured by any particular set of choices, the choice or value $\mathbf{s}_i \in \sigma_i$ for any particular σ_i is not constrained by the choice or value of any of the others.

   b. Information-theoretically, this implies that knowing the value or specific choice $\mathbf{s}_i \in \sigma_i$ for σ_i provides no information about the corresponding value or choice $\mathbf{s}_j \in \sigma_j$ for $\sigma_{j \neq i}$.

4. Call this "extensional" because it focuses solely on the legality of all possible sets of choices from the specified ranges of variation.
   a. Intensional versions will look at the mechanisms or properties that dictate how those choices are actually made (at how the variation actually occurs).
   b. The extensional version is simpler: it simply reduces the issue of whether or not the entities are independent to the question of whether any of the possible choice sets $\{\mathbf{s}_1, \mathbf{s}_2, \ldots \mathbf{s}_k\}$ violate condition **c**.
   c. If they do, then the set is not independent; if not, then it is.

B. Examples

1. Logic: some axiom is independent of another set of axioms (such as the axiom of choice being independent of the other axioms of set theory)
   a. Elements $\mathbf{x}_i$: axioms
   b. Space of choices σ_i: two-element set of truth-values {TRUE, FALSE}
   c. Condition **c**: resulting theory (with all the axioms) is consistent
2. Algebra: independent variables
   a. Elements $\mathbf{x}_i$: the variables themselves
   b. Space of choices σ_i: real numbers (say)
   c. Condition **c**: presumed equations true of those values
3. Probability: set of events
   a. Elements $\mathbf{x}_i$: events types
   b. Space of choices σ_i: probability of occurrence (being instantiated)
   c. Condition **c**: ?
      i. Note: the probability of joint occurrence of two independent events is the product of their individual probabilities because the probability of the second is *unaffected by the probability of the first*, in just the way required by the model: if the first happens, then the second is neither more nor less probable than if the first does not happen.
4. Geometry: length and width
   a. Elements $\mathbf{x}_i$: abstract rectangle parts (of a common rectangle)
   b. Space of choices σ_i: space of (measurable) extent—e.g., real numbers
   c. Condition **c**: governing constraints of Euclidean geometry

C. Comments on condition **c**

1. This condition is *extremely important*, in assessing independence claims
2. Consider rectangles
   a. If arbitrary, then **c** imposes no additional constraint
   b. But if we were to "change the subject"—and just consider rectangles of unit area, say—then the length and width would shift, to being *completely dependent*
   c. Or if surface of a sphere, then neither wholly independent, nor wholly dependent.

IV. Extensional constraint

A. Symmetry
   1. The notion of *independence* is often used symmetrically
   2. This fits with practice, where we talk about sets *as a whole*
      a. Probability: as above, two (or more) events are independent just in case the *probability of their joint occurrence is the product of the probabilities of their individual occurrence* (products are symmetric)Algebra: dependent or independent equations or variables
   3. But we can get an asymmetrical reading, by shifting attention to dependence

B. Extensional **dependence**
   1. Given above, call an entity **x** (**extensionally) dependent** on a set of others {$\mathbf{x}_1$, $\mathbf{x}_2$, … $\mathbf{x}_k$} just in case its choice or value $\mathbf{s} \in \sigma$ in a given situation {$\mathbf{s}_1 \in \sigma_1$, $\mathbf{s}_2 \in \sigma 2$, … $\mathbf{s}_i \in \sigma_i$}) is *fixed*.
   2. And (as above), call **x (extensionally) independent** of a set of others {$\mathbf{x}_1$, $\mathbf{x}_2$, … $\mathbf{x}_k$} just in case its choice or value $\mathbf{s} \in \sigma$ in a given situation {$\mathbf{s}_1 \in \sigma_1$, $\mathbf{s}_2 \in \sigma 2$, … $\mathbf{s}_i \in \sigma_i$}) is *unconstrained*.

C. **Extensional** *(partial)* **constraint**
   1. There's a great deal of room between (extensional) dependence & (extensional) independence. This is an extremely important fact.
   2. Both are extremal cases
      a. Independence: *no constraints at all*
      b. Dependence: *complete constraint*
   3. Whole region in the middle: of *partial constraint*
   4. Cf. figure 2 (assume **k**=2)
      a. Independence (of $\mathbf{x}_4$ of $\mathbf{x}_1$)
         i. Given a particular choice $\mathbf{s}_1 \in \sigma_1$ for $\mathbf{x}_1$, …
         ii. The choice for $\mathbf{x}_4$ remains *entirely open* (whole original set σ_4)
      b. Dependence (of $\mathbf{x}_2$ on $\mathbf{x}_1$)
         i. Given same choice $\mathbf{s}_1 \in \sigma_1$ for $\mathbf{x}_1$,…
         ii. The choice $\mathbf{s}_2 \in \sigma_2$ for $\mathbf{x}_2$ is *fixed*.
      c. Middle region of (partial) *constraint*
         i. Where choice $\mathbf{s}_1 \in \sigma_1$ for $\mathbf{x}_1$ means that the choices for $\mathbf{x}_3$ are *narrowed to a set* $\sigma_3' \subset \sigma_3$ *that nonetheless has more than a single element.*

D. Definitions
   1. **Partial constraint:** middle region, between dependence and independence
   2. **Constraint:** non-independence (dependence or partial constraint—i.e., left hand or middle regions of figure 2)
   3. E.g., say that **x** is *constrained by* **y** if either
      a. **x** *depends on* **y**, or
      b. **x** is *partially constrained by* **y**.

E. Consequences
   1. Non-independence (constraint): weaker than dependence
   2. Non-dependence: weaker than independence
   3. Symmetry:

a. Independence: symmetrical
b. Partial constraint: symmetrical
   i. $\mathbf{x}_i$ fails to be independent of $\mathbf{x}_{j\neq i}$ just in case the choice for $\mathbf{x}_j$ constrains the corresponding choice for $\mathbf{x}_i$ (which means that for at least one choice $\mathbf{s}_j \in \sigma_j$ for $\mathbf{x}_j$, there is a choice $\mathbf{s}_i \in \sigma_i$ for $\mathbf{x}_j$ such that $\{\mathbf{s}_i, \mathbf{s}_j\}$ violates c.
   ii. But this shows that the situation is symmetric.
   iii. For suppose that the choice for $\mathbf{x}_j$ were made first, and σ_i selected. That would preclude choosing σ_j for $\mathbf{x}_j$. So $\mathbf{x}_j$ is not independent of $\mathbf{x}_j$, either—even though neither of $\mathbf{x}_i$ or $\mathbf{x}_j$ need be dependent on the other, either.
c. Dependence: asymmetrical
   i. To say that $\mathbf{x}_1$ depends on $\mathbf{x}_2$ does not imply that $\mathbf{x}_2$ depends on $\mathbf{x}_1$
   ii. This is what you would expect

F. Converses (if $\mathbf{x}_i$ bears r to $\mathbf{x}_{j\neq i}$, what relation does $\mathbf{x}_{j\neq i}$, bear to $\mathbf{x}_i$?)
1. Dependence: *constraint (non-independence)*
   a. Not independence (as one might have thought)[2]
      i. But its opposite: *non-independence* (i.e., constraint: either partial constraint or more dependence)I.e., if $\mathbf{x}_1$ extensionally depends on $\mathbf{x}_2 \ldots \mathbf{x}_k$, then none of the $\mathbf{x}_2 \ldots \mathbf{x}_k$ can possibly be independent of $\mathbf{x}$; instead, they will all *partially depend* on, and thus be constrained, by $\mathbf{x}_1$.
   b. I.e., converse of left-hand side is middle and left-hand side together; not right-hand side.
   c. This is because of the extensional reading's underlying reliance on *pure correlation*, which is itself a symmetrical notion.
   d. Fits with mathematical or calculational account
      i. If $\mathbf{a} = \mathbf{b}^2$, then **a** depends on **b**, whereas **b** is *constrained* by **a** (because it can be either positive or negative).
      ii. No way that **b** can be claimed to be independent of **a**.
   e. For discussion, say that extensional dependence is only **weakly asymmetric**.
   f. Relationship of *strong* asymmetry, allowing $\mathbf{x}_2 \ldots \mathbf{x}_k$ to be *independent* of $\mathbf{x}_1$, even if $\mathbf{x}_1$ *depends* on $\mathbf{x}_2 \ldots \mathbf{x}_k$, will have to await an intensional reading.
2. Converse of constraint = constraint
   a. I.e., if **a** is constrained by **b** (partially or completely), **b** will similarly be constrained by **a**.
   b. So (this will matter to FSM): *extensional constraint is symmetrical*

G. Binary cases
1. If choice sets σ_i are binary—i.e., have exactly two elements—then the middle region collapses, and dependence and independence come out as exhaustive categories, with no region of partial constraint opening up between them.
2. Only two options
   a. Either $\mathbf{x}_2 \ldots \mathbf{x}_k$ restrict the choice $\mathbf{s}_1$ for $\mathbf{x}_1$ to a single element of σ_1, in which case $\mathbf{x}_1$ is dependent on $\mathbf{x}_2 \ldots \mathbf{x}_k$;

---

[2]Except in degenerate cases: Where $\mathbf{x}_1$ depends on $\mathbf{x}_2 \ldots \mathbf{x}_i$ because there is always only one possible choice $\mathbf{s}_1 \in \sigma_1$ — i.e., when $\sigma_1 = \{\mathbf{s}_1\}$ — or when the choice sets are binary, as discussed in a subsequent point.

   b. Or else $\mathbf{s}_1$ can be either of the two elements of σ_1, in which case $\mathbf{x}_1$ is independent of $\mathbf{x}_2 \ldots \mathbf{x}_k$.
3. Familiar, e.g. in logic, where choice sets are often taken to be the binary set of truth-values
   a. I.e., if the choices are binary, there is no such thing as *partial constraint*
   b. Either a set of claims dictates the truth of another one; or it does not.
4. Would be a mistake, though, to generalise from this special case.
5. As soon as questions grow more complex (even including *possible situations* in which a claim is true) dependence and independence will no longer, in general, be exhaustive categories.

V. Computation I: Semantics dependent on syntax

A. Intro
   1. How does this analysis of (extensional) dependence and independence apply to the situation of computating?
   2. Inversion
      a. FSM claims syntax (operations) independent of semantics
      b. But note: extensional independence is symmetrical
      c. So can ask inverse question: is semantics independent of syntax?

B. Scenario
   1. We will use a participatory example
      a. α designates β.
      b. α is causally dependent on β.
   2. Critic's argument
      a. Reinterpret α as designating $\gamma \neq \beta$;
      b. Observe that α's behaviour remains unchanged;
      c. Conclude: α is processed in virtue of effective connections, not semantic connections
      d. So FSM (behaviour independent of semantics) remains true.

C. Assessment
   1. Claim that the argument fails
   2. Crucial fact: local, causal interactions with semantic values that critic dismisses as "mere coöccurrence" are, in point of fact, *essential.*
   3. Relation between syntax and semantics, that is to say, is not just one of coincidental coöccurrence; it is one of *necessary overlap*.
   4. In general, that is
      a. The existence of an effective connection with the referent …
      b. Plays a role in establishing the fact that …
      c. That referent *is* the referent.
   5. That is, the coöccurrence critic, in making this move, is being a *quick change artist*!

D. Example: (as before) an internal cases: octal numerals
   1. Type '4+5' in to a compute
   2. It responds with '11'.
   3. Might think it had made a mistake.

4. Then write the following program:
   — for i := 1 to 20
     begin print i
     end
5. And suppose it prints out
   — 1 2 3 4 5 6 7 10 11 12 13 14 15 16 17 20
6. Similarly:
   a. 'length("collaboration")' ⇒ '15'
   b. Drag dozen files to the trash; it asks you to confirm that you want to delete '14' files
7. Etc.
8. Soon conclude: *it is using octal numerals*
9. What warrants this conclusion?
10. Coherence, writ large—including semantic realm, since counting is involved (in both directions)
11. Something like: coordination between direct causal interaction with, and representation of, a semantic realm.
12. Hard to argue: "no, in fact it is using base 10; it is just unremittingly confused." (Would be to act like Humpty Dumpty—taking the meaning of words into your own hands.)

E. Claim
1. When it involves direct causal interaction with the semantic domain, as well as coherent correspondence to that domain, the behaviour of a system cannot be ruled out as figuring in what warrants its having the interpretation that it does.
2. Note: this is a modest claim
   a. All this is necessary (n order for it to be true) is that the behaviour *figure* (i.e., play some kind of role) in establishing the interpretation
   b. It is not necessary (for it to be true) that behavior on its own be enough to *establish* it.
   c. So it is neutral with respect to origins of the interpretation (attributed or realist).

F. Examples
1. Embedded robots walking and driving around, constructing "maps" of the territory they explore, with respect to which they make decisions about where to plug in for the night.
2. Processes watching over the Internet, routing traffic around particularly busy nodes, and sending up-to-the-millisecond traffic reports to gateways scattered around the world.
3. Mail servers sitting on local area networks, opening and breaking connections to client processes in order to distribute, collect, and redistribute electronic mail.
4. Etc.

G. Status of this analysis
1. Properly speaking (as we said), we cannot establish claim that semantics depends on behaviour without a theory of semantics
2. But note:
   a. *No one has ever understood these systems* except under semantic interpretation,
   b. No one has ever interpreted them except with reference to their semantic domains—domains for which they are designed, in which they are deployed, and in which, once

deployed, they become effective, consequential, agents.

c. Cf. programmers' annoyance at Searle.
d. Contention: programmers don't think that "semantics is up for grabs"

3. Human reference
   a. Consider proposals made about the establishment of reference and interpretation in people (apart from computational considerations).
   b. They all make reference to some kind of effective engagement with the subject matter.
   c. Examples
      i. Causal theories of reference
      ii. Theories that take perception and action as essential (Harnad etc.)
      iii. Physical embodiment of the intentional agent
      iv. "Biological function" of Milikan et al.
      v. Dretske: causal efficacy of counter-factual correlation in learning
      vi. Fodor: asymmetrical dependence
      vii. Wittgenstein: "meaning as use"
      viii. Vygotsky: activity theory
      ix. Etc.
   d. All assume that *some form of coordinated interaction with the semantic realm*, not just a detached representational or deductive attitude towards it, plays a role in establishing the fact that that realm is the realm of semantic interpretation.
4. Anyone who wants to argue that the involvement of computers in their semantic realms is irrelevant to *their* semantic interpretation would thus have to take whatever is right about these accounts, and show that it *fails to be true of computers*
   a. This in spite of the fact that, as the wealth of examples in chapter 3 (the ones we went over in class last week) demonstrated, the range of kinds of participatory involvement of computational systems seems, on the face of it and on reflection as well, not to have any obvious or even natural limits beyond those enforced by the physical substrate, to which we humans are subject as well.
   b. Moreover, such an argument could not rely on any claim that computers are formal, or derivative, or syntactic, or the like—since those properties are exactly what is being contested.
   c. Note, too, that a number of diachronic properties—such as surviving and reproducing only because of standing in various "true" semantic relations to their subject matter-—are also manifestly true of computers in the wild;
   d. Not only is that why we build and distribute the ones we build, but even internal processes and threads, such as those generated by predictive algorithms and search routines, are selected and prosper because of their relative behavioural advantages.
   e. Finally, some of the elements of the semantic realms in question, such as specific threads, dead-locks, file headers, and the like, are things we humans know about almost entirely through interaction with computational *representations* of them, such as error messages, print-outs, etc.; the only systems that have direct effective engagement with these things are computers, not people.

5. In sum, INVOLVEMENT—fact that computers traffic in their own semantic realms, as illustrated in so many of the examples we have seen—makes it overwhelmingly likely that that *effective engagement will at least play a role in establishing their semantic interpretation.*
6. So, even lacking a theory of semantics, hardly contentious to suppose that semantics is not, in general, independent of behaviour — especially for a participatory system.

H. Conclusion
1. It is time to bring this all to a conclusion
2. Semantics on syntax:
   a. What we have argued is that, on an extensional reading of the issues of dependence and independence, **semantics is not independent of syntax**
   b. We have not, however, argued for the opposite: that semantics is *dependent* on syntax.
   c. Rather, we have suggested that the most likely situation is the middle reading: a relation of *partial constraint*
   d. That is: semantics is *partially constrained* by syntax (or behaviour)
3. Other direction
   a. The negative reading of formality, however, didn't say anything about whether *semantics was dependent on syntax*.
   b. Rather, the made it claim ran the other way around: it was about *syntax being independent of semantics*.
   c. So what do we have to say about that? What implication, for that thesis, is implicit in the results we've amassed so far?
   d. The answer is straightforward: because we are dealing with extensional versions, *the claims go the other way 'round as well.*
   e. Converse of *constraint* is *constraint*
   f. I.e., have both directions (extensionally)
      i. Syntax (operations, behaviour) *partially constrains* semantics (interpretation)
      ii. Semantics (interpretation) *partially constrains* syntax (behaviour)
4. In sum: FSM is *false if analysed extensionally*
5. Some will say: was never how it was supposed to be read
6. So: if we are to rescue the formality thesis, we will have to look at a *stronger notion of independence*.

VI. Intensional independence

A. Tune in next time …

—— *end of file* ——