Enough of metatheoretic abstractions.

In Part II of this essay, I have sought to survey basic neglected problems that confront any serious attempt to develop an integrated science of mental phenomena. The issues raised have had a disheveled diversity within which the only evident unifying theme is the extraordinary difficulty we can expect to encounter even when conjecturing *laws of cogitation, nevermind establishing them as laws, whose SLese quality is not an embarrassment. Indeed, it would be easy to conclude that mentation can have so little humanly fathomable systemacy that academics who fancy themselves as cognitive scientists should seek more honest employment elsewhere. My intent, however, has been not to put quietus to SLese reconstruction of folk psychology but to break ground for the foundations on which this must build if it is to achieve whatever may be its potential. So by rights, this essay should close by demonstrating how disciplined SLese thinking can deepen our understanding of some particular mental phenomenon of classic interest. Unhappily, the extent to which I can bring off that desideratum will elicit few sighs of gratification. Even so, it will be a useful exercise to review the issues abstractly examined earlier by seeing how they arise when one attempts to formulate principles under which commonsensical perceivings are causally responsive to the environmental events about which our percepts are putatively informative.

How do our perceptions relate to the external world? For any sentential clause 'that-p' expressive of a possible perceptual judgment, it is an important commonsense truism that

(54) For any observer o, if o perceives that-p, then p.

For example,

(54.1) For any observer o, if o perceives that-the-sun-is-shining, then the sun is shining.

As a scientific generality, however, (54) is severely defective in several instructive ways. First of all it entails, contrary to fact, that perception is always veridical. Once we make clear that 'o perceives that ...' is to be read here in its strictly psychological sense for which a philosopher might prefer 'It perceptually appears to o that ...', we can easily enough hang a layman's qualification on (54) by replacing 'then p' therein with 'then probably p'. But a technical science insists on more determinate appraisals of such probabilities; and more importantly, since the probability of p given o's perceiving that-p is strongly conditional on additional features of the perceiver's local circumstances, we would want to detail what these are and how perceptual accuracy is affected by them.

Secondly, although (54) generalizes over all objects in a vaguely specified domain ("observers") that would be clarified by our working out the determinants of perceptual accuracy, its second occurrence of 'p' is schematic not for a nominal but for a statement, and cannot meaningfully be quantified over.[47] As philosophers

---

[47] Although not all philosophers agree, I suggest that the single most important insight into ontology ever achieved by technical philosophy is the Quinian principle that only the nominal components of sentences can meaningfully be replaced by placeholders (logical variables) bound by quantifiers. Thus, a sentence of form 'F(a)' logically entails 'There is something such that F(it)' only if the occurrence of term or phrase 'a' in this context is purporting to name (designate, refer to) something. To illustrate, from the sentence 'My pencil is dull' wherein 'my pencil' is a nominal but 'dull' is an adjective, I can properly infer 'There is something that is dull' but not 'There is something that my pencil is'. But if I paraphrase(???) this premise as 'My pencil has (the property) dullness' by nominalizing its predicate, then 'There is something that my pencil has' follows by impeccable logic. Similarly, I cannot meaningfully infer 'John sneezed because something' or 'There is something that John sneezed because' from 'John sneezed because his nose tickled', anymore than I can infer 'John sneezed or something' from 'John sneezed or Mary blinked'. But I can validly infer 'John's sneezing was due to something' from 'John's sneezing was due to his nose's tickling.' To non-philosophers, this distinction may well seem stupifyingly recondite; yet unless it is heeded with deep respect, philosophies of language, logic, and ontology are at grave risk of incoherence.

---

of language since Tarski well know, replacing any such scheme for semantic use/mention interchange with an assertable generality relating concepts to external reality is a task of enormous difficulty. (Note that although conversion of 'then p' in (54)

to 'then that-p is true' allows quantification over 'that-p' as placeholder for the name of a proposition, the resulting generality no longer links percepts with their external objects.) Making articulate what nameable entities in the external world, abstract or concrete, stand in lawful relations to what nameable constituents of perceivings is a major goal for any serious science of perception.

Thirdly, (54.1) illustrates that even were perceivings always veridical, few ordinary-language sentences would instantiate 'p' acceptably in (54). For when John-today and Mary-yesterday both see that-the-sun-is-shining, the perceived external event which (54.1) tries to identify by 'the sun is shining' is evidently not the same for John-today as it is for Mary-yesterday. That is, (54.1) is shorthand for something like

(54.1a) For any observer o, if o perceives that-the-sun-is-shining, then probably the sun is shining in o's vicinity.

But (54.1a) is not of form (54); so what then is the perceptual principle, schema or generality, this embodies? Other than abandoning (54) altogether, we have just two options: (a) We can retain (54) unaltered by requiring each substitution for 'p'

therein to be a sentence 'p_i' that signifies just one determinate state of affairs unaccomodated to any particular o in (54)'s domain while the perceptual content that-p_i expressed by 'p_i' is free of demonstrative elements that allow the object of an o's perceiving-that-p_i to depend on o. Otherwise, (b) we must put aside (54)'s simplistic use/mention interchange in favor of some

(55) For any observer o, if o perceives that-p, then probably #[that-p,o] ,

wherein '#[that-p,o]' is schematic for the statement produced from 'o' and any acceptable instantiation of 'p' by a still-to-be-devised algorithm # which, for any observer-name 'o_j' and any relative clause 'that-p_i' adequately expressive of a possible perceptual judgment, transforms <'that-p_i','o_j'> into a sentence whose gerundization would designate what o_j's truthful perceiving that-p_i would be a perception of. Thus, # must enable (55) to capture the form of

(55.1) For any observer o, if o perceives that-the-sun-is-shining-here-and-now, then probably the sun is shining in o's vicinity.

(Note that (55.1) differs from (54.1a) in making explicit certain demonstrative components of perceptual content that are presumably implicit in (54.1a).) Unlike prospect (a), which is a non-starter, seeking to cash out # in meta-schema (55) would be a reasonable and indeed valuable enterprise for the psycho-philosophy of perception. But to bring this off with any success we need to articulate perceptual contents far more richly than provided by ordinary-language verbalizations of percepts, and to develop some sophistication in the SLese details of how perceptual events arise.

Finally, even if (54) or better (55) were impeccable in all other respects, it would only schematize diagnostic laws lacking any explanatory force. For it is clear even to folk psychology that the conditionality of (54) is not that of schema

(56) For any cognizer o, o's ϕing-that-p probably brings it about that p.

Some Psi-verbs, notably 'endeavors', do indeed yield plausible instantiations of

(56) for suitably chosen 'p'; but in all likelihood (56) is never true even roughly for any 'p' when 'ϕ' is 'perceive'. The closest counterpart of (55) whose conditionality is commonsensically causal is

(57) For any observer o, if #[that-p,o], a probable result is that o perceives that-p,

except that whereas the probability in (55) is merely suspect, in (57) it is vanishingly small. Schema (57) illustrates what we want of an explanatory theory of distal perception (or rather, it is an unSLesed precursor of that, as (4) is of (9) in Chapter 1); but it makes plain our need to put flesh on connection schema # (more technically, to spell out the details of locus structure in laws of perceptual arousal by distal macro-stimuli) and calls even more loudly for expansion of (57)'s if-clause to include conditions with which state of affairs #[that-p,o] must be supplemented if o's perceiving that-p is to be produced with respectably high probability.

Of course, perceptual generalities needn't be causal in order to have scientific merit. But that is required if we are to understand why people perceive as they do with what accuracy under what circumstances; and our notes on the illusory simplicity of (54) leave little reason to hope that instructive acausal world/percept covariations will be any easier to come by. Let us consider, therefore, how development of the causal story might commence.

How are percepts differentiated?

Turn over this book and, after removing its dust jacket if still there, inspect its binding. What do you see, and how? Commonsense is quick to answer that what you see is the entity your percept is of, namely the current temporal stage of this book or more precisely some state of affairs in which this book-stage is a constituent, while your activated perceptual representation of that is the means by which you see it. But here we want a deeper reading of these questions: What you perceive is to be conveyed by a statement of form 'I see Q' wherein 'Q' details

the character of this seeing (i.e. its content) by a relative clause distinct from what would describe your perceiving's character had this book's binding been shaped or pigmented differently. And <u>how</u> you see is to be answered in terms of the mechanism through which the bookish state of your near environment gives rise to your seeing Q**ly** rather than Q*ly for any SLese contrast Q* to content alternative Q.

Here are some of the relative clauses that folk psychology would favor for surmising how you may have reacted perceptually in this experiment:

(58-1)		that this book is blue,
(58-2)		that this book's binding is blue,
(58-3)		that this book's binding is middling-dark blue,
(58-4)		that this book has a blue binding,
(58-5)	___ sees	that this blue thing is rectangular,
(58-6)		that this rectangular thing is blue,
(58-7)		that this thing is rectangularly blue,
(58-8)		that this binding has printing only on its spine,
(58-9)		that one's hand partly covers this book,
(58-10)		that one's hand has ragged nails,

(where the idiom of self-report would replace the pronoun in (58-9,10) by 'my'), and so on for enormously many additions to this list. Right off, then, we have a major problem in explaining your commonplace "information pick-up": Precisely what <u>is</u> the percept to be accounted for here? Our uncertainty about one particular mental event which probably not even you can elevate to the status of observational datum for our epistemic community is not the issue; rather, it is which of these ordinary-language perceptual prospects, if any, are reasonable to posit about some book-viewer's perceptual experience in order to inquire how they might lawfully have arisen. Presuming for the moment that at least one of predicates (58) became true of you, did incompatibilities thwart your satisfying more than one of them? Or were several true of you simultaneously and, if so, were these co-occurrences merely coincidental or did some analytically necessitate others? Above all, are one or two of these, or certain others that belong on the list, perceptually <u>primary</u> in that we pretty well have to work out the theory of such primary perceivings before we can get leverage on percepts that are in various ways derivative from these?

Consider, for example, (58-2) vs. (58-3). Commonsensically, you can doubt, or hope, or expect that-a-is-blue for some suitable nominal 'a' without doubting/hoping/expecting that a is any more determinate shade of blue. But can you see that-a-is-blue without perforce seeing also that-a-is-$ish-blue for some shade qualifier '$'? Let us provisionally agree that (58-2) is a-derivative from (58-3) in essentially the way that weighing-roughly-128-lbs. is an abstraction from (inter alia) weighing-127.3852-lbs. But if so, (58-3) presumably holds for you only by rounding off, in turn, an even more determinate coloration percept which no English phrase adequately conveys. Might the latter's specificity then also analytically include whatever it takes to make (58-8) true of you as well? (I will later argue not, but a case can be made either way.) And in similar vein, is (58-1) merely an abstraction from if not elliptic for (58-2)? Or, alternatively, did (58-2)'s holding for you leave open or even interfere with (58-1)'s holding as well?

Although (58-1) can be construed simply as shorthand for (58-2), its strict reading suggests a perceptual-content difference illustrated more explicitly by the variation within (58-2,4) and (58-5,6,7). Commonsense disputes that seeing this book either as blue or as having-a-blue-binding is identical with seeing as blue just this book's binding. Nor does ordinary language regard (58-5,6,7) as paraphrastically equivalent. But are some of these merely derivative from others--e.g., might (58-5) and (58-6) be entailed by (58-7), and (58-2) by (58-4)--or are the properties these respectively represent so distinct that competetion may prevent co-occurrence of more than one in each group? (We shall return to this comparison later.)

That we find queries such as these perplexing makes plain that the ordinary-language relative clauses you might spontaneously use to tell others what you perceived when looking at this book's binding, or to which you might assent if asked whether such-and-so is what you saw, appear dubiously adequate to express the distinctive character of your perceiving with the precision wanted for a target of scientific explanation. We could scarcely expect otherwise; for our perceivings precede our

verbal reports thereof by a causal gap assuredly large enough to preclude tight correspondence between seeing and saying even under the most favorable conditions of self-report. That is, although from your voicing that you saw that-this-book's-binding-is-blue we can fairly infer that your just-preceding percept was probably one which, conjoint with inter alia your current language habits and motivation, is apt to elicit some verbalization in a class of rough synonyms for 'this book's binding is blue' (or 'this ... looks blue' or 'this ... seems blue'), in all likelihood there are a great many other perceptual contents that could also have prompted you to this same verbalization and are diagnosed by the latter with scarcely less plausibility, if not more, than the percept you actually had. To be sure, technical research on perception often waives self-report in favor of more sensitive non-verbal indicators of input reception such as stimulus matching/ordering and discrimination thresholds; but mere inaccuracy of self-reports is not our point at issue here. Rather, if the phrases afforded by ordinary language for differentiating percepts are not even roughly in one-one (or many-one) correspondence with the perceptual distinctions that seem needed to account for perception-mediated behavior, by what linguistic devices <u>are</u> these distinctions to be drawn by a science of perception? For example, instead of merely inspecting this book's binding, you might try to sort a large number of variously shaded blue chips into a spatial layout whose between-chip spacings correspond to the degrees of color similarity you see among them. Arguably, in order to make these comparisons you have to see each chip as a distinct shade of blue. If so, how are we to individuate these percept-shadings in conjectured accounts of their lawful evocation in you when ordinary English does not give us the words to do so? Nevermind how we might learn for sure what the to-be-accounted-for events in fact are; the deep problem is how can we even conceive what they distinctively might be in the first place.

Were our deficiencies in perceptual predicates merely suboptimal precision of distinctions already roughed in by our extant language, their alleviation would prima facie be largely routine. (Technical science has had several centuries

of practice at concept refinements.) For example, we can easily imagine adding labels to our language, and training ourselves to use them discriminatingly, for all the different colors displayed, say, in the Munsell Color Atlas.[48] But surely

---

[48]It is by no means certain that when your perceiving of this book should be described by a predicate of form '___ sees that this book's binding is $ish in color', the more finely I can myself discriminate and verbalize shades of color, the more closely I can approximate the specificity of your percept by some substitution for '$' from my expertise vocabulary. But even were this to be so, there is rather more to establishing that vocabulary as one our epistemic community can use than just publishing a carefully ordered and indexed collection of color chips. Access to such a physical atlas immediately gives me the use of color-comparison predicates of form '___ sees that this thing is the same color as the atlas chip labeled $L_{\$}$'; but that is still some distance from my acquiring the ability to make meaningful use of a predicate '___ sees that this thing is $ish in color' in which '$ish' is some adjectival variant of label $L_{\$}$. Arguably--though some profound obscurities in the nature of language-as-we-use-it troubles this thesis--my use of this latter locution is not "meaningful" in the fashion wanted unless I can use color-qualifier '$' in nonrelational color judgments (e.g., 'This thing appears $ish to me') that are highly predictive of the color comparisons I might then make between things I judge to $ish/non-$ish and chips in the atlas. Specifically, if I judge object a by itself to be $ish, the chip to which I then match a upon inspection of the atlas should be the one labeled $L_{\$}$.

---

not all possible perceivings, e.g. some by isolated aborigines, or chimpanzees, or human infants, or pigeons, are even roughly synonymous with any expressions in our shared adult language either now or in foreseeable future enrichments thereof. And if so, in what terms can we speculate about the possible character of those percepts? As a baseline for discussion, I give you

Posit. A condition of internal arousal is not a percept, or at least is not identifiable by us as one, unless we can give it an individuating description comprising a verb (notably 'perceives' or one we take to demark a particular style of perceiving) followed by a phrase formalizable as 'that-$\underline{F}(\alpha,\beta)$', with singular '$\underline{F}(\_,\_)$' therein the schema of some class of English subject/predicate sentences. More specifically, 'α' and 'β' are to be tuples 'α' = $\langle$'α_1',...,'α_m'$\rangle$ ($m \geq 1$) and 'β' = $\langle$'β_1',...,'β_n'$\rangle$ ($n \geq 0$) of possibly-complex symbols such that each 'β_j' can be tagged by some English expression '$\underline{b}_j$' (perhaps only a word-radical), and each 'α_i' by some English nominal '$\underline{a}_i$', for which (writing '$\underline{a}$' for $\langle$'$\underline{a}_1$',

..., '$\underline{a}_m$'> and '$\underline{b}$' for <'$\underline{b}_1$',..., '$\underline{b}_n$'>) '$\underline{F}(\_,\underline{b})$' is a well-formed English predicate including quantifiers not / and '$\underline{F}(\underline{a},\underline{b})$' an assertable English sentence. (Predicate '$\underline{F}(\_,\underline{b})$' is of course pulled out of '$\underline{F}(\alpha,\beta)$' by replacing the latter's 'α'-terms by placeholders and its 'β'-symbols by the corresponding elements of '$\underline{b}$'. There is no imputation here that 'α' has any similarity of meaning to '$\underline{a}$', or 'β' to '$\underline{b}$'. Rather, these English "tags" serve to associate each symbol 'α_i' or 'β_j' with a grammatical type of expression exemplified by its tag.)

Evidently this Posit needs commentary. First, it allows that the English tag associated with 'α_i' or 'β_j' might be just 'α_i' or 'β_j' itself. So it subsumes cases where '$\underline{F}(\_,\beta)$' is already an English predicate or '$\underline{F}(\alpha,\beta)$' itself an English sentence. Secondly, there is reason for the Posit's elaborate wording; for although I would have preferred it to declare simply that a percept's description is to have the structure of a singular (i.e. unquantified) subject/predicate proposition, the force of saying that is unhelpfully obscure. Even so, whatever propositionally structured mental contents may be,* they are something that everyday English tries to characterize by declarative sentences converted to relative clauses; and the Posit stipulates that a percept recognizable as such must be describable by putting symbolic elements into an English sentence-frame in such fashion that if each of these elements <u>were</u> to be added to English as a meaningful expression of its tagged grammatical type, then this percept-description <u>would</u> be distinguished by an English relative clause with the grammatical structure of the English sentence-frame it now embeds. Predicate '___ perceives that-$\underline{F}(\alpha,\beta)$' is a theoretical construct whose meaning for us is defined by the psychonomic theory we make with it; and to complete our contention that it describes a percept with content that-$\underline{F}(\alpha,\beta)$, despite '$\underline{F}(\alpha,\beta)$' not being a sentence in our own language, we must conjecture *laws governing this predicate wherein its 'α' and 'β' constituents play roles rather similar to what commonsense takes to be the distinctive contributions of '$\underline{a}$' and '$\underline{b}$' to verbalized conjectured causes and effects of perceivings-that-$\underline{F}(\underline{a},\underline{b})$ when '$\underline{F}(\underline{a},\underline{b})$' is a sentence in English.

But isn't the Posit's preoccupation with English intolerably chauvinistic? Not really; or at least not insofar as sentence-frames that might be accepted under the Posit's counterpart in another language have English translations. Admittedly, the possibility does remain that perceptual structure has variants not adequately captured by the syntax of any communication system currently recognized as linguistic. But we have little hope of imagining what those might be until we come to understand the constitution (i.e., a/t-derivational nature) of properties we now take to be paradigmatically perceptual.

On the other hand, even within the framework of English sentences, may not the Posit be overly narrow in limiting percepts to singular subject/predicate form? Major issues arise here. The umbrella question is, given a set $\{p_i\}$ of propositions logically interconnected in ways to be illustrated, is it possible for these all to be contents of an observer's near-simultaneous perceivings and, if so, do some of these perceivings-that-p_i necessarily derive either causally or abstractively from others? Let us consider some cases couched in ordinary English.

Suppose that your inspection of this book's binding made some version of (58-9) true of you, say '___ sees that this$_1$ covers this$_2$'. (Idiom would say '___ sees that this covers that', but I prefer 'that' to remain univocal here.) The sentence that expresses this percept's content is a paradigm of subject/predicate composition in ascribing a predicate of nearly minimal grammatical complexity ('__ covers __') to a pair of nominals ('this$_1$','this$_2$') whose own internal syntax is likewise minimal. Now: Is it possible that you also saw <u>that</u>-<u>something</u>-<u>covers</u>-<u>this$_2$</u>, or <u>that</u>-<u>this$_1$</u>-<u>covers</u>-<u>something</u>, or <u>that</u>-<u>something</u>-<u>covers</u>-<u>something</u>-<u>else</u>? Modern logic formalizes the sentences describing these latter three content possibilities as '$(\exists \underline{x})\underline{P}(\underline{x},\underline{a}_2)$', '$(\exists \underline{y})\underline{P}(\underline{a}_1,\underline{y})$', and '$(\exists \underline{x},\underline{y})\underline{P}(\underline{x},\underline{y})$', respectively, all of which are different one-way logical consequences of the sentence '$\underline{P}(\underline{a}_1,\underline{a}_2)$' formalizing the content of seeing <u>that</u>-<u>this$_1$</u>-<u>covers</u>-<u>this$_2$</u>. So be sure to understand these existential generalizations to give seeing that-something-covers-this$_2$ (etc.) a content somewhat

different from that of seeing that-this$_1$-covers-this$_2$. (Although idiom finds it easy to use 'something' as a bare demonstrative in perceptual contexts, its sense as a quantifier in, e.g., 'Not everything is covered by something' is what we want here.) But can you genuinely perceive that-something-covers-this$_2$ as distinct from perceiving some particular thing as covering this$_2$? Clearly a fully particularized perceiving can convince you of the generality that-something-covers-this$_2$; but is this existentially quantified awareness too a perceiving or is it instead only a post-perceptual belief? We have already Posited against the former, but that may only call the Posit into question. For three obscurities have now become obtrusive: First, if you can see that-something-covers-this$_2$, can you do so except as a result of some perceiving whose content is a fully particularized that-α_1-covers-this$_2$? Secondly, if the former requires the latter, is its derivation a causal production or an analytic abstraction? And finally, if this derivation is causal, can both its antecedent and its consequent be perceivings?

Whether one perceiving can causally evoke another may at first seem to be largely a matter of definition. For if we stipulate, reasonably enough, that perception is the first phase of input processing to which our commonsense language of intentionality (Psi-verb talk) applies, we might then also prefer that any cogitations aroused in turn by percepts are to be classified as post-perceptual ideation. However, the foregone certainty that perceptual processes are dynamically auto-regressive--i.e., that whatever we take to be the perceptual state of observer $\underline{s}$ at time $\underline{t}$ is a major source (at the appropriate level of molar causality) of $\underline{s}$'s perceivings at time $\underline{t}+\Delta$--pretty well requires that we allow percepts to be prevailingly caused in part by other percepts. Still, we want to distinguish hard-core perceivings from their less-sensuous cogitive consequences even though this contrast

surely spreads over a graded series of differences on which the commonsense division between perceiving and (more intellectualized) thinking/believing would be arbitrary were it not so broadly vague. So it seems appropriate to envision multiple stages of perceptual processing whose details will emerge only as we work out the nature of cogitive contents but which are ordered (perhaps only partially ordered) by a normal sequence of activations from peripheral input passing through distinguishable steps of perceiving over into central nonsensuous thought. If so, we can waive qualms about the generic admissibility of generalized propositions as perceptual contents in favor of doubt only that these can be contents of perception's earliest stages. In keeping with that move, we modify the term 'percept' in the Posit's opening clause as 'first-stage percept' or 'primary percept'.

Alternatively, the Posit's restriction to (singular) subject/predicate content structures can be dropped simply by replacing '($\underline{m} \geq 1$)' therein by '($\underline{m} \geq 0$)'. However, that would defeat half of the Posit's dual purpose, which is not merely to sketch how we can get a conceptual handle on perceptual contents beyond the reach of ordinary English, but also to urge that in perception there is something especially basic about singular subject/predicate composition.

Let us accept, then, that you can see that-something-covers-this$_2$, or that-something-covers-something-else, as a result of seeing that-this$_1$-covers-this$_2$. But resulting how? Could all or most of the latter be abstractively contained in the former, more or less as seeing that-this-is-darkish-blue presumably abstracts into seeing that-this-is-blue? The answer is neither clear to me nor is really needed for present purposes except insofar as the emended Posit would be easier to defend were perceiving that-$(\exists \underline{x})\underline{P}(\underline{x},\underline{a}_2)$ to be a-derivative from perceiving that-$\underline{P}(\underline{a}_1,\underline{a}_2)$ rather than caused by it. Even so, the issue is basic for perceptual theory, as other examples can bring out more forcefully.

Suppose that instead of merely turning this book over, you laid it down, withdrew your hand, and became aware that-not-everything-is-covered-by-something.

Considering the complex pure generality of its content, which modern logic would parse by formalism '$\sim(\forall \underline{y})(\exists \underline{x})\underline{P}(\underline{x},\underline{y})$', could this awareness possibly be a *percept*? Maybe, maybe not; introspection seems indecisive. But however it is to be classified, we can reasonably presume it to be due to your awareness *that-nothing-covers-this*$_2$, which surely is as sensuously vivid as any percept suggested by list (58). And although the latter's content formalization as '$\sim(\exists \underline{x})\underline{P}(\underline{x},\underline{a}_2)$' or '$(\forall \underline{x})\sim\underline{P}(\underline{x},\underline{a}_2)$' again exhibits the structure of a complex generality which seems far more suited to post-perceptual ideation than for perceptual immediacy, commonsense would be outraged by insistance that you can't really perceive a thing's lack of encumberances. Must we concede, then, that your seeing that-nothing-covers-this$_2$ is a first-stage percept which refutes even the amended Posit? Not if we can argue that your seeing-that-$\sim(\exists \underline{x})\underline{P}(\underline{x},\underline{a}_2)$ derives, either causally or by abstraction, from some other perceiving of yours with fully singular content. And to my own introspective sensitivities that does indeed seem correct. Contrary to what a logician might prefer, you assuredly do not perceive that-nothing-covers-this$_2$ by inductive inference from an array of perceivings $\{$that-α_i-does-not-cover-this$_2\}$. Rather, what does seem psychonomically plausible is that your no-covering awareness arises from a first-stage percept whose content is some *that-this*$_2$*-is-*β*ish* with a predicate component β rather similar to the meaning of '___ is unobstructed' except for being conceptually elemental, *not* built up from still-more primitive predicates by negation and/or quantification. Moreover, the meaning gap between *that-this*$_2$*-is-*β*ish* and *that-nothing-covers-this*$_2$, which does encorporate negation and quantification, appears sufficiently large that the latter cannot reasonably be viewed as just an analytic abstraction from the former. If so, your passing from seeing that-this$_2$-is-βish to seeing-that-nothing-covers-this$_2$ is a causal progression by some nomic principle which neither embodies any entailment schema recognized by modern logic nor always produces the latter as concommitant to the former, as would be required were the one to be analytically contained in the other.

Similarly, if introspection is correct in allowing you to see that-nothing-covers-this$_2$ without necessarily seeing also that-not-everything-is-covered-by-something, or to see that-this$_1$-covers-this$_2$ at times bereft of seeing that-something-covers-this$_2$ or that-something-covers-something-else, then these existentially generalized seeings can only be all or in part causal consequences of, not wholly abstractions from, the fully singular seeings that give rise to them. Even so, there still remains some question whether the seemingly chancy accompaniment of seeing-that-$\underline{P}(\underline{a}_1,\underline{a}_2)$ by seeing-that-$(\exists\underline{x})\underline{P}(\underline{x},\underline{a}_2)$ is the latter's being indeed a stage of perceptual arousal distinct from the former, or whether this separation might instead be just an illusion of unreliable verbal reporting. That is, with brutal oversimplification, perhaps your seeing-that-something-covers-this$_2$ consists of your saying-to-yourself the words, 'Something covers this', in response to seeing-that-this$_1$-covers-this$_2$ without causal mediation by any existentially generalized perceiving.

Roughly speaking, the more syntactically and/or conceptually complex is the relative clause by which we describe a putative percept, the farther downstream in a process sequence this perceiving seems likely to be from any primitive onset of perception. Consider, for example, (58-8). Even apart from the negation and quantification discernable therein, the concepts of <u>printing</u> and <u>book-spine</u> within its content implicate a high degree of "interpretation"--i.e., integration of present experience with an intricate residue of past cogitations--whose evocation is surely consequent upon some less intellectualized perceivings of pigmentations. Conjecturably, these primary perceivings are an ensemble having descriptions suggested by {'___sees that this$_i$ thing-part is $\beta_i$ish in color'} and {'___ sees that this$_j$ thing-part contrasts β_{jk}-wise with this$_k$ one'}. Or should we conjecture instead that your first-stage perceiving here is describable by just one '___ sees that this thing is β*ish in color', wherein the β*-qualifier, though not a color concept available in English, manages somehow to encode all the color

features that you can attribute to this book in later stages of perceptual and post-perceptual cogitation? The latter does not preclude your also having an array of derivative perceptual properties {seeing that-this$_1$-thing-part-is-$\beta_1$ish-in-color, etc.; but whether these are then analytic abstractions from a single seeing that-this-thing-is-β*ish-in-color, are separate causal consequences of it, or are some mixture thereof, remains an achingly open question.

The ensemble of simultaneous perceivings just envisioned points toward still another instructive puzzle case in how percepts are to be descriptively/ontologically distinguished. In simplification of (58-9,10), suppose it is possible for you to (a) see that-this$_1$-covers-this$_2$ while simultaneously seeing also that-this$_1$-is-ragged. How does this dual perceiving differ from (b) seeing that-this$_1$-covers-this$_2$-and-this$_1$-is-ragged, and the latter from (c) seeing that-this$_1$-covers-this$_2$-and-is-ragged? One answer, which seems clearly wrong to me, is that these are simply three different ways to describe the very same perceptual content, or--a slightly weaker claim--that they make nearly but not quite identical abstractions from a base percept which ordinary English cannot clearly distinguish from them. Whether (a) might differ from (b) is tested by considering whether both halves of the following biconditional are true:

(59-1) For any propositions that-p and that-q, if any observer o perceives that-p-and-q, then o perceives that-p and perceives that-q.

(59-2) For any propositions that-p and that-q, if any observer o perceives that-p and perceives that-q, then o perceives that-p-and-q.

(59-1) can easily be defended on grounds that any perceiving whose content is a conjunction of propositions analytically contains perceiving each one of them. But (59-2) implies that there is nothing more to propositional conjunction than just co-occurrence, which is sufficiently implausible to discourage equating (a) with (b) even though we seldom heed this distinction in everyday perception-talk. Even more conspicuous is the manifest difference between (b) and (c), at least if the syntax

of our descriptions thereof is to be taken seriously. For the content of (b) attributes a 3-place predicate to a 3-component subject albeit giving two of those subject components a special content similarity, whereas the content of (c) is a binary predication. If this descriptive difference is not just a quirk of English idiom but reflects a genuine distinction in perceptual organization, it becomes incumbent on perception theory to explain precisely how these differ, both in psychonomic function and in constitution. In all likelihood, (a) is simply an abstraction from a variety of more determinate perceivings, (b) and (c) among others. But given that neither (b) nor (c) abstracts from the other, do these normally occur in successive stages of a causal progression (and if so, which comes first?), are they on separate causal paths which can nevertheless run off synchronically, or does arousal of the one competetively preclude conjoint activation of the other?

Concommitant to but deeper than such functional questions about (b) vs. (c) is the issue of how these differ constitutionally: Could that be largely variance in the locus structure of micro-events from which (b) and (c) respectively abstract, comparable to the locus-structural difference between uttering, hearing, or reading the word-string (b') 'This covers that and this is ragged' in contrast to the string (c') 'This covers that and is ragged'? The core of your emitting/receiving (b') or (c') is a temporal sequence of motoric/sensory verbal subpatterns; and the only peripheral difference between these two utterings/hearings/readings (neither of which manifests much locus-structural similarity to the sensory input that produces your (b)-seeing or (c)-seeing) is that sequence (b') repeats a subpattern which (c') contains only once. So given that the abstraction base of your molar (b)-seeing or (c)-seeing must have a t-core (cf. p. 150, above) comprising an ensemble of micro-events distributed in space and time, do the nominals in this percept's English description correspond to particular subsets of percept-constituting micro-events in such fashion that in case (b), but not (c), there are three so-named micro-arrays of which two share a certain attributive character? And is the micro-molar reason why you cannot effectively utter/hear/read (b') and (c') simultaneously

(despite there being ways in which this is physically possible) a close counterpart of why (b) and (c) are in large measure competetive? We shall explore this prospect, that the constitutional nature of having a percept is basically of a kind with receiving/emitting a sentence, at some length following a pause to take stock of what we have been doing and where this will lead.

Whence and whither.

Our ultimate aim in this chapter is to consider how far a science of perception may possibly go in establishing well-SLesed laws under which distinctive features of an observer's external surround elicit one determinate perceiving rather than another. But the very first step in that undertaking is working out a language within which we can characterize the assorted perceptual alternatives whose dependencies upon environmental elicitors are to be disclosed. (We cannot intelligibly say why someone perceives $\underline{Q}_i$ly rather than $\underline{Q}_j$ly until we can replace '$\underline{Q}_i$ly' and '$\underline{Q}_j$ly' by words that literally mean something.)

Seeking to verbalize specific to-be-accounted-for percepts (perceptual states of mind) throws us immediately into a love/hate affair with ordinary language. For not only does ordinary language contain unboundedly many predicates that are commonsensically perceptual, technical psychology has no present grounds on which to classify any psychonomic property $\underline{P}$ as "perceptual" except by arguing that $\underline{P}$-ness appears to be the sort of thing that commonsense perception talk is about. Yet ordinary language quickly proves inadequate for a technical science of perception. For on one hand we quickly find ourselves wanting to distinguish perceptual alternatives far more articulately than ordinary language avails, even while, on the other hand, it is far from clear which variations in ordinary-language perceptual locutions reflect significant differences in reality rather than stylistic artifacts of ellipsis, paraphrase, metaphor, allusion, and other context-dependent fluidities of phrasing that are the bane of a hard science.

So how do we escape this bind? Our answer is to hypothesize provisionally that were everyday English to be expanded into a capacity to express all possible descriptive concepts of the various grammatical types we now recognize, to be enriched without limit in all ways that our working vocabularies do in fact increase with maturation and experience, then use of this perfected language according to the rules by which we now construct English perceptual predicates would provide individuating identifications of all percept alternatives that our science thereof wants to study. Then for inquiry into perceptual issues that cut across the details of particular instances, we can take our examples from everyday English with all the intuitional/folk-psychological leverage that gives us on the question at hand, while expecting that whatever we conjecture or provisionally conclude from such ordinary-language cases should apply as well to percepts described with greater technical adequacy whenever we become able to bring that off. Meanwhile, as we develop theories about the psychonomic nature of commonsense percepts, we shall find ourselves regimenting, restricting, and perhaps eventually even modifying the grammar of ordinary-language perception talk--in short, the sort of bootstraps lift by which any technical science moves beyond its commonsense origin. But to commence, we must tug on whatever lacings we initially find in place.

Although ordinary English allows sentence radical 'o perceives ___' to be completed by expressions of widely diverse grammatical types, I have stipulated (with cogent but scarcely irresistable argument in Chapter 5) that only completions of form 'that-p', with 'p' a grammatically well-formed declarative sentence, are to be accepted for now as describing mental states of the perceptual sort. But that still gives us an enormous field of purported perceptual properties signified by predicates of form '___ perceives that-p', some of which, if possible of realization at all, are surely derivative one way or another from more basic percepts which are the most seemly candidates for governance by laws of distal perception. To focus this concern, let $\{$'___ sees that p_i'$\}$ comprise the predicates formed from all the sentences $\{$'p_i'$\}$

whose only descriptive terms, or grammatical variants thereof, are contained in the relative clause of some commonsensically straightforward perceptual predicate '___ sees that p_0'. Given that the latter is true of some o, which of the former must then also be true of o, which must be false of o, and most importantly, what understanding of perceptual organization do we acquire from attempting to decipher why these entailments and exclusions hold?

Our reflections on commonsensical perceptual possibilities (58) have noted three kinds of content interplay which any serious theory of perception needs to recognize and give some account of. One is the connection between seeing-that-p_0 and seeing-that-p_i when sentence 'p_i' derives from 'p_0' by converting some predicative concept therein into a weakened, looser, less determinate version of that distinguisher, e.g., rarefying 'middling-dark-blue' into just 'blue'. In this case, given that there indeed exists a seeing-that-p_0 property signified by predicate '___ sees that p_0', we have three primary prospects for the semantic status of '___ sees that p_i' when commonsense understands this to be entailed by '___ sees that p_0'. One is to say that '___ sees that p_i' lacks the precision required for a predicate to represent anything, so that strictly speaking there is no such property as seeing-that-p_i. Another is taking '___ sees that p_i' to signify loosely on certain occasions of its usage the very same property that '___ sees that p_0' signifies more precisely, i.e., to hold that when sentence 'o sees that p_0' truthfully represents o as seeing-that-p_0, the predicate in 'o sees that p_i' also represents seeing-that-p_0. And finally, we can allow that although seeing-that-p_0 and seeing-that-p_i are distinct properties, the latter is abstractively contained in the former, just as Rectangularity is an analytic abstraction from (inter alia) Squareness. For reasons that need not be aired here, the first two alternatives are to be shunned if at all possible. In contrast, despite ontological qualms that Option 3 may occasion, treating this rarefied seeing-that-p_i as an abstraction from the more determinate seeing-that-p_0 is simply one more application of the a-derivational thinking whose SLese formalisms

have been exercised repeatedly in preceding chapters and without which it scarcely seems possible to do technical science. That some percepts are analytically contained in or "supervenient" upon others has no outset import for perception theory beyond preparing us to recognize that this is one salient way in which a to-be-explained perceiving may be due to another. But it is a nice question whether some levels of perceptual abstraction are not better behaved in causal regularities than are others and, if so, by what signs (e.g. description features) can we pick out percepts on these psychonomically preferred levels? Similarly but more directly germane to our hopes of finding principles under which perceptual events are governed, it seems exceedingly unlikely that there exist law-schemata or meta-laws (cf. p. 215 above) that subsume perceptual variables with indifference to their abstraction levels. It follows that any conjectured *principle of perception worth taking seriously must be carefully restricted to, inter alia, some particular abstraction level for which we have worked out specifications.

(Note, however, that even when '___ sees that p_i' is a linguistic attenuation of '___ sees that p_0' as just described, it remains conceivable that we sometimes understand these predicates to designate separable percepts such that, on the one hand, seeing-that-p_0 not only does not necessitate but may even interfere with simultaneous seeing-that-p_i, while conversely, seeing-that-p_i needs not be accompanied by any seeing-that-p_0^* characterized by a content sentence 'p_0^*' of which 'p_i' is an attenuation. How this case differs from the one wherein seeing-that-p_i is an abstraction from seeing-that-p_0 is one of the clarifications we expect from any serious account of what percepts are.)

Secondly, we have encountered the question of how seeing-that-p_i relates to seeing-that-p_0 when sentence 'p_i' is formed from a singular (fully determinate) subject/predicate sentence 'p_0' by quantifying over nominals in the latter. Although it is conceivable that seeing-that-p_i in this case is abstractively contained in seeing-that-p_0 when all quantification in 'p_i' is existential--later I will defend

a thesis not far from that--introspection urges not merely that there is in general a real separation here, but also that when quantified percepts arise they are causal consequences of fully singular ones, the latter being what distal stimulation evokes most directly. Even if that is not always true, singular percepts are by far the easier to model, and confront us most starkly with the irreducible subject/predicate core of perceptual structure--which is why we shall say little more here about quantified contents. Even so, making place for an account of this apparent difference in kind between quantified percepts and singular ones is a condition of adequacy on any model of perception.

Finally, we have noted that the content clauses of two singular perceptual predicates can contain exactly the same concept elements, yet assign these to different roles in syntactic frames that may themselves differ. Whether such grammatical variations reflect genuine contrasts in the percepts so described, and if so what is their nature, cuts to the heart of the fundamental psychonomic problem of mentality: When we attribute to some o a percept (or any other moded thought) with sententially characterized content, what does that say about o beyond o's mere thinking a cluster of ideas corresponding to the list of this sentence's morphemes? That is, what in a propositionally structured percept is more than just the aggregate of its content elements, and how does that something-more affect the dynamics of thinking? This question has remained so profoundly ignored in cognitive science that verbalizing any constitutional models of content structure, even simplistic ones of dubious merit, cannot help but significantly advance our comprehension of this matter.

There is still another important perception-theoretic obscurity visible in our preceding examples. It is plain from these that ordinary-language perceptual predicates are pervaded by demonstrative terms. That is no great matter if these merely token the inadequacy of extant English for describing percepts with the precision and detail that an unboundedly enriched English would allow. But the prospect that demonstratives may not be eliminable from percept descriptions even in principle threatens cognitive science with shipwreck. For if '$\underline{P}(\_)$' is a predicate,

perceptual or otherwise, that contains a demonstrative whose linguistic force is relative on each occasion of this predicate's usage to some aspect of that local circumstance, can this context-dependent locution '$\underline{P}(\_)$' then possibly participate in any law-statement $\underline{L}$ that not only generalizes over an open domain but also aspires to convey the same representation of how the world works on all occasions of $\underline{L}$'s contemplation by the epistemic community concerned with this generality? Although the answer, I fear, is that No, this is not possible, we can still seek ways for a science of perception to admit perceptual demonstratives even while evading condemnation under this conclusion. (Indeed, we shall later make considerable effort to accomplish that.) But clearly, some coming to terms with demonstratives must be given high priority by any serious concern for communicable perceptual regularities.

We shall now undertake deeper probing of the issues just scanned by observing their differential portrayals in two strongly contrastive models of how perceptual properties might be embodied in complexes of brain conditions. (Strictly speaking, localization of percepts in the brain is a minor detail easily waived by these models; but I include it because we have every reason to feel sure that the t-cores of commonsensically conceived mental events are indeed molar occurrences within nervous systems.) One model cashes out the widespread notion that thinking is somehow an internalized exercising of language; the other reflects the view that perception is basically imagistic. Both models are skeletal in the sense that each leaves many major details unspecified. It is, however, of considerable value to appreciate how perception theory can pursue one or another forceful direction of model development without premature commitment to particulars that are best left open until their choice becomes well motivated.

No firm conclusions will emerge from these model comparisons, for although I shall eventually urge that one of the two can largely be dismissed, it would be foolish to suggest that the other wins by default over more sophisticated accounts

not yet aired. But models need to develop motivated complexity by evolving under criticism; and this chapter's aspiration is not final resolution of any perception-theoretic uncertainties, but layout of foundation issues brought into clear focus by some outrageously idealized models capable of goading us to conceive of superior alternatives. Meanwhile, in an immediate application, we shall later see how usefully even a primitive constitutional model can guide our sorting of perceptual properties into variables, as required for us to get on with search for functional regularities in perception.

What percepts might be: Two models.

The prospect of instructive constitutional parallel, between perceivings-that-p and peripheral occurrences of acoustic or graphic English sentences whose meaning is that-p, suggests a model for the nature of percepts valuable not so much for its likely accuracy as for the fix it gives us on what is possible. According to this, the inner-sentence model of perception, the t-core locus of o's perceiving that-p is an array $\langle o^*_a, o^*_b, \ldots \rangle$ of disjoint brain-part stages in o--i.e., each o^*_i is a restricted though perhaps discontinuous region of o's neural tissue during some particular portion of the time-interval spanned by this perceiving--which contains the proposition that-p as follows: First, for each i = a,b,..., either the totality of o^*_i's activation state or a certain abstraction from this is the fragment of perceptual content, i.e. concept, expressed by some meaning-subdivision in a sentence 'p' which in our language asserts that-p (or which would assert that-p for us were our vocabulary to be suitably enriched). We can allow that o^*_i's activation state embodies as many different perceptual concepts simultaneously as there are different abstractions from that state picked out by our language for percept-description. (Thus if o^*_i's state embodies the middling-dark-blue concept, it also embodies the more abstract blue.) And secondly, the grammar-conveyed organization which makes the proposition expressed by sentence 'p' much more than just an an aggregate of meanings elicited piecemeal by morphemes ordered arbitrarily in 'p' is some configuration of locus-structural

properties/relations on o's brain-stage regions $\langle o_a^*, o_b^*, \ldots \rangle$.

Crudely illustrated, the inner-sentence proposal is that your seeing that-this_1-covers-this_2 takes place in three regions o_a^*, o_b^*, o_c^* of your brain during the time of this seeing (which regions may or may not be simply three disjoint stages of the same continuant neural register) such that o_a^* sensuously thinks the this_1-concept, o_b^* and o_c^* sensuously think covers and this_2, respectively, and these three disjoint local thoughts are woven into a propositional attribution of covers to $\langle \text{this}_1, \text{this}_2 \rangle$ by some special compound $F(\_,\_,\_)$ of properties and relations on $\langle o_a^*, o_b^*, o_c^* \rangle$ that constitutes the syntactic frame of a binary predication. The ingredients of structure F might include aspects of these brain-regions' temporal sequence, their geometric layout, their sizes and synaptic connections, even perhaps transient conditions independent of their concept-embodying activation-state features in counterpart to grammatical inflictions in a spoken sentence. But ontological details of F's composition remain an outstanding mystery.

[The inner-sentence model posits a basic distinction between (a) brain-region characteristics that, in rough correspondence/analogy to the spatio-temporal layout of morpheme sites in sentences of an uninflected language, constitute the syntactic framework of propositional thoughts, and (b) their activation states embodying concept elements. And I have further presumed that inner-syntax conditions (a), or their abstraction bases, are mainly of the sort repeatedly referred to in previous chapters as "locus structure." I have deliberately evaded specifics on what that comprises, for theories of causality can as yet offer little more than open speculations about it. But it paradigmatically comprises those conditions γ, notably space-time displacements or whatever else may constitute excursive preconditions of causal connection, which appear as domain constraints on the loci of events governed by laws written in the t-core detail formalized by (9'), p. 35 above. So allowing "locus structure" to include monadic attributes as well as polyadic relations, we can say that a causal system's locus-structural ingredients are in essence just the properties (including relations) that figure in its laws' domain preconditions once all t-derivational constructions have

been written out of its variables.

[[Unhappily, since the metaprinciple of Domain Constriction allows properties that are values of one law's variables to be domain preconditions of another, this does not provide nearly so clear a division as we might wish between brain regions' locus-structural properties on one hand and their activation-state features on the other. Even so, all that an abstract inner-sentence model requires of a structure/activation distinction is for the multifarious characteristics of brain regions to partition into two rather different sorts, one serving to embody the syntax of propositional thought while the other embodies elements of conceptual content. Conjecturably, most of the former should appear as domain preconditions in laws governing t-core mental events while the latter are mostly values of those laws' process variables. But we need no strong presumptions about that at this time, except that to acknowledge what I suggest is a fundamental ontological difference between a causal law's domain preconditions and properties governed by production principles within that domain, we should provisionally stipulate that properties which count as "structural" are not to be analytically dependent on any conditions we treat as "activational."]]

Although the inner-sentence model undoubtedly bears some nontrivial resemblance to the realities of perceptual constitution, I offer its present sketch as no more than a simplistic heurism which needs at best extensive elaboration and more likely major modifications. In particular, we have not considered how logical connectives and quantifiers are to be encorporated in inner-sentences of supra-minimal complexity. Even so, not only does it clarify what might be the nature of propositionally structured percepts, and from there give ready access to specific ways in which different percepts may be linked both causally and abstractively, its contrast with the yet-to-be-described "inner-picture" model of perceiving valuably illuminates the most fundamental character of percepts and other thoughts, namely, their subject/predicate articulation.

My previously-argued thesis that peripheral-to-central percept sequences always commence with fully singular subject/predicate contents may well be overly extreme.

Yet even the most logically complex perceivings are tokened by their relative-clause descriptions to embed a subject/predicate structure regardless whether their subject-slots are filled by determinate nominals or only by quantified placeholders. By all rights, then, our key to unlocking the nature of perception should be found in the functional/constitutional difference between subject and predicate in perceivings of minimal syntactic complexity. So for intimations of generality as well as brevity, let us shorten (58-5,6,7) to

(58-5') ___ sees that this B-thing is R ,

(58-6') ___ sees that this R-thing is B ,

(58-7') ___ sees that this thing is Rly B ,

and contemplate their divergence.

The leading question about these percept descriptions is whether their grammatical variance is not just linguistic paraphrase but reflects genuine contrasts in perceptual content, as intuition urges. All three contain exactly the same conceptual ingredients, but profess different apportionments of them between subject and predicate. In the inner-sentence model of these perceivings, that distinction is both real and perspicuous: For an observer-stage o of which (58-5') or (58-6') or (58-7') is true, o's brain contains two regions o_a^* and o_b^* such that the activation state of o_a^* is, or abstracts into, the sensuous ideational content this-B-thing or this-R-thing or just this-thing respectively (where 'this' tokens additional content particulars which our descriptions of these contents cannot easily specify); the excitation state of o_b^* is, or abstractively includes, the sensuous R-idea or the B-idea or the Rly-B-idea, respectively; and certain structural relations between o_a^* and o_b^* establish whatever may be o_b^*'s content as predicated of whatever content is in structural subject-position o_a^*. Presuming that such syntax relations, even if to some extent transient, come about more or less independently of content activations in the brain regions so related, it is clear that establishing $\langle o_a^*, o_b^* \rangle$

as a subject/predicate frame allows the propositional content it contains in principle to attribute any concept within the range of o_b^*'s state alternatives to any concept in the state range of o_a^*. This includes the prospect of o_a^*-state/o_b^*-state feature sharing illustrated by

(60-1) ___ sees that this R-thing is Rly B ,

(60-2) ___ sees that this Rly-B-thing is B .

(In these, as in (58-5',6',7'), 'thing' is best viewed as a contentless syntax marker of subject-position.)

[If the Rly-B concept is a coordinate construction equivalent to R-and-B, as in rectangularly-blue, the inner-sentence model allows the within-subject structure of (60-2), or the within-predicate structure of (60-1) and (58-7'), to be captured by further partition of region o_i^* (i = a or b) into two locus-structurally distinguished subregions $o_i^* = \langle o_{i_1}^*, o_{i_2}^* \rangle$, with the states of $o_{i_1}^*$ and $o_{i_2}^*$ respectively embodying the R-idea and B-idea, or more determinately a this-R-idea and this-B-idea. (Depending on just what is in content surplus this, $o_{i_1}^*$ and $o_{i_2}^*$ may further partition into sub-subregions whose respective states split apart various facets of this-R and this-B.) For such compounds, seeing-that-α-is-Rly-B analytically contains seeing-that-α-is-B (or seeing-that-this-Rly-B-thing-is-β contains seeing-that-this-B-thing-is-β) by part/whole inclusion wherein the first is embodied in a certain ensemble of structural relations and nonrelational activation states of $\langle o_a^*, o_{b_1}^*, o_{b_2}^* \rangle$ whose subarray just for $\langle o_a^*, o_{b_2}^* \rangle$ constitutes the second. In contrast, if R in Rly-B is an "intensifier" of B, as in middling-dark-blue or (to illustrate demonstrative qualifiers) in this-blue, the inner-sentence model embodies both seeing-that-this-α-is-Rly-B and seeing that-α-is-B in the joint subject-state/predicate-state of the same region-pair $\langle o_a^*, o_b^* \rangle$, but with the content of o_b^* containing the Rly-B concept at one level of abstraction and the B concept at a higher one.

[[The inner-sentence model also cheerfully accepts that all of (58-5',6',7') might hold simultaneously for observer o, perhaps even in multiple embodiments. For if o contains several brain-region pairs $\langle o^*_{aj}, o^*_{bj}\rangle$ ($j = 1,2,\ldots$) with each pair $\langle o^*_{aj}, o^*_{bj}\rangle$ satisfying the structural conditions for it to be a subject/predicate frame, then similarity or difference among the subject-content states of $\{o^*_{aj}: j = 1,2,\ldots\}$, and among the predicate-content states of $\{o^*_{bj}: j = 1,2,\ldots\}$ is constrained only by nomic covariation, not constitutional overlap. Hence in principle, given a suitable configuration of causal antecedents (which, however, might be quite difficult to bring about), $\langle o^*_{a1}, o^*_{b1}\rangle$ and $\langle o^*_{a2}, o^*_{b2}\rangle$ might both contain that-this-R-thing-is-B, $\langle o^*_{a3}, o^*_{b3}\rangle$ might contain that-this-B-thing-is-R, and $\langle o^*_{a4}, o^*_{b4}\rangle$ might contain that-this-R-thing-is-nonB or even that-this-R-thing-is-nonR, with this-R-thing being the very same nominal concept embodied in the states of these various subject locations within macro-observer o.]]

The inner-picture account of perception, on the other hand, tells a very different story about (58-5',6',7'). This is the model that more or less identifies percepts with stimulus-driven images, although I shall leave it for you to judge how closely inner pictures in my sense resemble what has been the focus of recent controversy (cf. Kosslyn, 1980; Pylyshyn, 1981) on the nature of imagery. According to this model, seeing that-α-is-βish consists of o's having a brain region o* which is itself the subject-component in this perceptual proposition and of which the β-concept is predicated by o*'s activation state being, or abstractively embodying, content β. What gives o* its particular nominal character (as distinct from what this perceiving predicates of it) is some array of structural properties identified so far as we are able by the 'α'-locution (e.g., 'this-R-thing' vs. 'this-B-thing' vs. 'this-Rly-B-thing' for our present study cases, but also terms such as 'I', 'you', 'Mary', and 'the Smyths' in commonsense applications) in our description of this percept's subject. Just what is to count as a "structural" property, as

distinct from activation-atate features, remains widely open (cf. p. 235f., above) for later explication in whatever directions the model finds most congenial. But those structural properties of a perceiving's t-core locus that have been most explicit, or nearly so, in traditional intuitions about inner pictures (see Kosslyn, 1980, pp. 32-35, 131-134) are its geometric shape, size, and position in space-time. (Technically, the psychonomic shapes/sizes/positions of brain regions are undoubtedly best defined mainly in terms of neuronal interconnections and humoral/electrical capacities rather than physical space-time coordinates; but for the present overview, physical geometry is most heuristic.)

In the inner-picture model, if the sensuous R-idea and B-idea are, or abstract from, brain-region shape and activation state, respectively, the t-core of observer o's seeing that-this-R-thing-is-B would be some region o* of o's brain-stage having an R-type shape while undergoing B-featured activation. (If you read R as rectangular and B as blue, it should be intuitively clear how a brain region of a certain shape that encodes rectangularity, and throbbing with the pattern of activity standardly elicited by blue stimuli, can be viewed as depicting an external rectangular object's being blue in color. Be clear, however, that this structural property embodying the rectangular-idea is not required to be a shape at all, much less one with right-angled corners.) But if R is structural while B is activational, it is then impossible for any of o's perceivings to have a propositional content wherein the R-idea occurs predicatively or which includes the B-idea in the percept's subject-component--which is to say that (58-6') and (58-7') in this case are either unrealizable or are misleading paraphrases for (58-5'). Alternatively, if the inner-picture model posits that both the R-idea and the B-idea abstract from mutually independent dimensions of neural activity in the same brain region, o*'s total activation state might embody Rly-B while the nominal concept of which Rly-B is predicated by o*'s having this state may be no more than a bare demonstrative this-ness embodied, say, in region o_k^*'s space-time position. (More on this shortly.) But that still prevents

(58-5',6',7') from being distinguishable perceivings in the inner-picture model, for now (58-5') and (58-6') become misleading paraphrases for the syntactically correct (58-7'). And (60-1,2), read literally, become impossible in either case.

The inner/sentence/inner-picture divergence is summarized with enhancements in Table 1/ (pp. 241a,b). Yet that scarcely touches the ramifications of percepts being constituted in the one format rather than in the other. These are best unfolded by contrasting how these two models are disposed to answer certain large questions of which I have been flashing glimpses but have not yet laid out foursquare. To sharpen the bite of these questions (which apply not just to percepts but to cogitations of all modal persuasions), it is useful to acknowledge the commonsense essence of thought, namely, its "intentionality" or representational character. The epistemic job of our concepts is to be about other things and from there to form propositional compounds which, in some though not all modes of entertainment, comprise our "information" about selected aspects of the outside world. I have already argued (pp. 137ff., 140f., above) that our understanding of representational aboutness is still far too dim for this notion to be admissible in any basic *law of mentation. But competing models of perception can be valuably illuminated by their contrastive implications regarding what, in an observer o's external surround, could reasonably be represented, under some yet-to-come explication of aboutness, by what aspects of o's total perceptual condition.

Pursuant to our focus on the fully singular propositions which I have urged are prevailingly if not exclusively the contents of primary percepts, we can safely presume that whatever is represented by o's truthful seeing that-α-is-βish must be some event, a's-having-B, such that o's α-concept and β-concept respectively stand for (designate, signify, refer to) object a and property B.[49] (Given these

[49] The representation of properties by predicate-concepts is semantic-theoretically very tricky, owing first of all to the murky ontology of properties (see fn. 15, p. 99 above) and secondly to the failure of predicates to function grammatically like nominals--which is why I prefer to say that predicates "signify" rather than "refer." I trust that you will not begrudge me a certain initial glibness in this matter which will to some extent be shed as we proceed.

TABLE 1

IDEALIZED INNER-PICTURE VS. INNER-SENTENCE MODELS OF A PERCEPTUAL EVENT, o's-seeing-that-α-is-βish.

Common Premises:

C1. The t-core of o's-seeing-that-α-is-βish is a molar event, o*'s-having-$P_{\alpha\beta}$, in which o* is a possibly-scattered region of observer o's brain-stage, and $P_{\alpha\beta}$ is some complex pattern-property of o*. The subregions $\{o_j^*\}$ in any mereological partition of o* are in all likelihood distributed in time as well as in space.

C2. The properties (including relations) of o* and its subregions $\{o_j^*\}$ are of two disjoint kinds, structural and activational, at all levels of molar abstraction. We leave open the substantive nature of this difference except for the understanding that any molar property is activational if its abstraction base non-vacuously includes any micro-property that is activational. (Thus the micro-conditions upon which activity patterns supervene standardly include structural properties, whereas conversely, structural properties do not supervene even in part upon activations.)

Nonperceptual example. Suppose that o* is a pigmented surface patch partitionable as a disjoint array $\{o_j^*\}$ of roughly-square subpatches. Then paradigmatically, the shapes, sizes, physical locations, and inter-patch distances of these o_j^* are micro-structural properties which abstract into the shape, size, and location of o*; the "activational" properties of o*'s subpatches $\{o_j^*\}$ include their respective local pigmentations; and the degree to which o* as a whole is checkered in pigmentation--a nonmentalistic counterpart of perceptual pattern $P_{\alpha\beta}$--is a molar activational property of o* which supervenes upon both the pigmentations and the structural features of $\{o_j^*\}$. (For details, see Chapter 5, p. 162f.)

The Inner-Sentence Paradigm:

S1. The t-core locus of o's-seeing-that-α-is-βish is a pair of subregions, $o^* = \langle o_a^*, o_b^* \rangle$, such that $F(o_a^*, o_b^*)$ for some special complex structural condition $F(\_,\_)$ on pairs of brain-regions. That is, F is a compound of structural properties which includes at least one (anti-symmetric) relation. Structural condition $F(\_,\_)$ is the psychonomic embodiment of subject/predicate form, the two open positions therein providing for insertion of subject-content and predicate-content, respectively.

[Table 1, continued]

S2. The subject-content and predicate-content in o's-seeing-that-α-is-βish are certain activity patterns α and β, respectively, in o*'s F-demarked subject-location o_a^* and predicate-location o_b^*.

S3. $P_{\alpha\beta}$, the propositional content in o's-seeing-that-α-is-βish, has the relational composition $P_{\alpha\beta}(x,y) = F(x,y) \,\&\, \alpha(x) \,\&\, \beta(y)$. ('x' and 'y' here are logical placeholders.) Hence the t-core of o's-seeing-that-α-is-βish has composition $P_{\alpha\beta}(o^*) = F(o_a^*, o_b^*) \,\&\, \alpha(o_a^*) \,\&\, \beta(o_b^*)$.

S4. The proposition in o's-seeing-that-α-is-βish is just $P_{\alpha\beta}$ itself, i.e., is the same as this perceiving's propositional content.

S5. Logical complexity within the α-concept or β-concept is similarly embodied by structural relations and local activational properties of disjoint subregions of o_a^* or o_b^*, respectively.

The Inner-Picture Paradigm:

P1. The subject-content in t-core o*'s-having-$P_{\alpha\beta}$ of o's-seeing-that-α-is-βish is a compound $F_\alpha(\_)$ of o*'s structural properties. In principle F_α includes specification of o*'s location in space/time or some neural-connection counterpart thereof, and in practice is expected to do so.

P2. The predicate-content in o's-seeing-that-α-is-βish is an activity pattern β of o* as a whole.

P3. $P_{\alpha\beta}$, the propositional content in o's-seeing-that-α-is-βish, has the conjunctive composition $P_{\alpha\beta}(x) = F_\alpha(x) \,\&\, \beta(x)$. Hence this perceiving's t-core has composition $P_{\alpha\beta}(o^*) = F_\alpha(o^*) \,\&\, \beta(o^*)$.

P4. The proposition in o's seeing-that-α-is-βish is the event o*'s-having-$P_{\alpha\beta}$, i.e., o*'s-having-both-F_α-and-β. [The depictive model of propositions has two other main variants. One, adopted in the text for simplicity, holds that the proposition here is simply o*'s-having-β but with the α-concept included therein by virtue of F_α being in the "nature" or "essence" of o*. The other takes pattern $F_\alpha(\_) \,\&\, \beta(\_)$ to be itself the proposition, in prima facie agreement with S4 of inner-sentencing.]

P5. The sub-events from which o*'s-having-$P_{\alpha\beta}$ abstracts are also t-cores of other perceivings by o, as developed in the *Principle of Dense Depiction, p. 245 below.

designations by α and β, we can forego presumption of truth by saying only that this percept represents object a as having property B.) So there are three sorts of sub-propositional representation to be provided for by a model of perception: that of properties by predicate-concepts; that of objects by nominal-concepts; and finally--the ultimate challenge of propositional "structure"--representation by internal syntax of the compositional nexus that integrates compound entities. By "compositional nexus" I mean above all the Exemplification tie of objects to their attributes, and Co-exemplification of two or more attributes in a common bearer of them as distinct from their looser co-presence at different locations in a common scene. But other important instances are the connection between a molar object and its mereological parts (e.g., the inclusion of John in John-and-Mary), and an attribute's embodiment of its higher abstracta (e.g., containment of rectangularity in squareness.) With these points in mind, let us see how the inner-sentence and inner-picture models compare in their views on how densely, and with what sort of segregation, an observer's simultaneous perceivings represent his surround; what limitations there may be on what can be perceptually represented either predicatively or nominally; whether the objective world's compositional nexus is literally reproduced or is more flexibly represented in perception; and finally, in a complex intertwining of semantic issues, the extent to which perceptual demonstratives may be radically particular.

Representation of compositional connection.

What I have labeled "inner-sentence" and "inner-picture" models of perceiving are, of course, frameworks open to considerable arbitration in detail. But what I take to be generically definitive of inner-picturing is its auto-representation of (inter alia) Exemplification, Co-exemplification, and Part/Whole Inclusion, in contrast to what, for want of a better word, may be called the "extrinsic" representation of these integrative couplings by inner-sentences. Regarding Exemplification, inner-picture models take o's seeing that-α-is-βish to be a depiction of some

a's-having-B through o's containing a brain-region o* whose position, shape, size, or other still-unknown structural features somehow pick out object a as the external referent of o*, while o*'s activation state embodies a certain abstract character β which, under some still-obscure principle of predicate signification, stands for external property B. That is, the event a's-having-B is here depicted by the event o*'s-having-β, with the observer's brain-region o* itself being the nominal (i.e. α-concept) in this percept's propositional content even though there remains a story to tell about how the reference-fixing structural properties of o* figure in our description of this percept and not merely fix reference but do so in part by representing certain structural features of a. And if this percept's predicate-component is a conjunction, $\beta_1$ish-and-$\beta_2$ish, of concepts β_1 and β_2 that respectively signify properties B_1 and B_2, what depicts a's-having-B_1-and-B_2 is o*'s-having-both-β_1-and-β_2. That is, co-exemplification of B_1 and B_2 is here represented by the co-presence of patterns β_1 and β_2 in the total activation state of the same brain site.

[[In a variant of the inner-picture model to which summary statement P4 in Table 1 gives lead billing, we can say that when a's-being-B is perceptually depicted by o's brain-region o*'s having activation character β, what we refer to by the nominal phrase 'α' when describing this percept's content as that-α-is-βish is not o* in itself as a bare ontological particular but its structural condition F_α that selects a as the object this percept is of. So construed, the percept's subject-content is like its subject-content in that both are prima facie attributes albeit of different kinds that are inflexibly nominal and predicative, respectively; representation of a as having B now becomes co-exemplification (by o*) of structural features F_α and activation-state abstraction β; and the α-concept can be viewed as not just referring to a but also representing certain structural properties of a--albeit how F_α represents those may well be rather different from the manner in which β represents B. It is unclear whether these two variants of depictive repre-

sentation differ in anything beyond their manner of speaking.[49a] In any case, they agree in representing objective co-exemplification of attributes by co-exemplification of brain-regional activity patterns that respectively signify them.]]

[49a]The most important technicality in this is a stupifyingly recondite ontological puzzle: When $\underline{o}^*$ has structural condition $\underline{F}_\alpha$, is the latter just an accidental (contra essential) possession of $\underline{o}^*$, or does $\underline{F}_\alpha$ inhere in $\underline{o}^*$'s being the particular object it is even to the point of $\underline{o}^*$'s being virtually identical with $\underline{F}_\alpha$? To be specific, suppose that $\underline{L}(\underline{o}^*)$ is $\underline{o}^*$'s complete space-time location (i.e. what is specified by the totality of position coordinates for points within $\underline{o}^*$), while $\underline{o}^*$ is whatever we refer to by some descriptor such as 'The neurone-stage from which we got the first micro-electrode reading in subject No. 3 yesterday'. Then does $\underline{o}^*$'s-having-location-$\underline{L}(\underline{o}^*)$ consist in some substantival inhabitant of container space-time, namely $\underline{o}^*$, having $\underline{L}(\underline{o}^*)$ as a predicable accident, or might it not be instead that the subject of predication here *is* location $\underline{L}(\underline{o}^*)$--i.e., perhaps $\underline{L}(\underline{o}^*)$ and $\underline{o}^*$ are one and the same--while our nominal 'The neurone-stage from which ...' designates $\underline{L}(\underline{o}^*)$ by citing an individuating collection of qualities at that location. Whatever the ontological truth in this matter, it seems highly dubious that $\underline{o}^*$'s-being-at-$\underline{L}(\underline{o}^*)$ is an event of the sort that arise as effects in causal processes. We should look for this to figure in nomic conditionals not as a production but only as a domain precondition--despite the proclivity of classical physics to treat spatial location, split off from temporal position, as a dependent variable.

In contrast to depiction, the inner-sentence model conjectures that representation of $\underline{a}$'s-having-$\underline{B}$ in $\underline{o}$'s-seeing-that-α-is-βish consists of $\underline{o}$'s containing a brain-region pair $\langle \underline{o}^*_a, \underline{o}^*_b \rangle$ having structural features--especially relational ones--that make them a subject/predicate frame while the α-concept and β-concept are embodied in the separate activation states of $\underline{o}^*_a$ and $\underline{o}^*_b$, the nominal neither exemplifying the predicate nor being co-exemplified with it. Rather, the outer-world Exemplification in which object $\underline{a}$ stands to property $\underline{B}$ is represented in this model by some quite different relation, one *not* also involved in the $\underline{a}$'s-having-$\underline{B}$ event, that holds between one brain region whose α-patterned activity refers to $\underline{a}$ and some other brain region whose β-patterned activity signifies $\underline{B}$. And if this β-predicate is a conjunction β_1-*and*-β_2*ish*, the inner-sentence model partitions $\underline{o}^*_b$ as two disjoint subregions, $\underline{o}^*_{b1}$ and $\underline{o}^*_{b2}$, whose activation states respectively embody the β_1 and β_2 concepts while some structural relation between $\underline{o}^*_{b1}$ and $\underline{o}^*_{b2}$ (which needn't be more than $\underline{o}^*_{b1}$ and $\underline{o}^*_{b2}$ each being linked with a common $\underline{o}^*_a$ in whatever fashion constitutes an inner-sentence subject/predicate frame) demarks this

subframe as a conjunctive predication. Note that β_1 and β_2 here could well be the complete activation states of $\underline{o}^*_{b1}$ and $\underline{o}^*_{b2}$, respectively, whereas this conjunctive predication's depictive construal requires β_1 and β_2 to be noncompetitive proper abstractions from the complete activation state of a common $\underline{o}^*$.

As for Part/Whole Inclusion, which figures importantly in the *Principle of Dense Depiction immediately below, the inner-picture model takes object-$\underline{a}_1$'s-being-part-of-object-$\underline{a}_2$ to be perceptually represented by a pair $\langle \underline{o}^*_1, \underline{o}^*_2 \rangle$ of brain regions such that $\underline{o}^*_1$ is physically a subregion of $\underline{o}^*_2$ while each $\underline{o}^*_i$ ($\underline{i}$ = 1,2) has the structural features needed to make $\underline{a}_i$ the referent of $\underline{o}^*_i$ in an ordinary subject/predicate depiction. In contrast, the inner-sentence paradigm of $\underline{o}$'s seeing that-α_1-is-part-of-α_2 would be $\underline{o}$'s containing a brain-region triple $\langle \underline{o}^*_{a_1}, \underline{o}^*_{a_2}, \underline{o}^*_b \rangle$ whose structural layout establishes $\underline{o}^*_{a_1}$ and $\underline{o}^*_{a_2}$ as subject-positions to which $\underline{o}^*_b$ is attached as a binary predicate-position, while the activation states of $\underline{o}^*_{a_1}$, $\underline{o}^*_{a_2}$, and $\underline{o}^*_b$ respectively embody the α_1-concept, α_2-concept, and part-of concept which in turn respectively stand for the external $\underline{a}_1$-object, $\underline{a}_2$-object, and mereological-inclusion relation. With appropriate adjustments of $\underline{o}^*_b$'s activation state, the same format holds for inner-sentence representation of any other external relation between objects $\underline{a}_1$ and $\underline{a}_2$.

Perceptual density and segregation.

Whenever an observer perceives that-α-is-βish (where α may be an $\underline{n}$-tuple $\langle \alpha_1, \ldots, \alpha_n \rangle$ and β an $\underline{n}$-adic relational predicate), any $\underline{a}$'s-having-$\underline{B}$ event represented by this percept is necessarily accompanied by many others which may be called its "factive concommitants." Since we do not here require a careful account of this notion, I shall declare somewhat arbitrarily that the factive concommitants of $\underline{a}$'s-being-$\underline{B}$ comprise all $\underline{a}_i$'s-having-$\underline{B}_j$ events wherein $\underline{a}_i$ is either $\underline{a}$ itself or is a mereological part of $\underline{a}$. (I would prefer factive concommitance to concentrate on events that are supervenient on the same array of external micro-events from which $\underline{a}$'s-having-$\underline{B}$ abstracts; but that restriction is hard to pin down.) Then the question

of perceptual "density" concerns the extent to which, when o truthfully perceives that-α-is-βish, the perceivable factive concommitants of the event this percept represents are also represented in o's synchronic totality of percepts. (We disregard factive concommitants that seem beyond the reach of perceptual representation in real life, such as quantum-mechanistic states of a's individual atoms.) And perceptual "segregation" concerns the extent to which the arrays of neural microevents which respectively constitute o's perceptual representations of different factive concommitants of the same a's-having-B are disjoint.

Intuitively, inner pictures are abundant representations in that each part of an inner picture represents some corresponding part of the larger event represented by the picture as a whole (cf. Kosslyn, 1980, p. 33). More specifically,

*Principle of Dense Depiction [PDD]. If a's-having-B is perceptually represented in observer o by an inner picture wherein o's brain-region o^*_k refers to a by virtue of o^*_k's structural properties, then for any brain-region o^*_{ki} that is a physical part of o^*_k: (1) o^*_{ki} has structural properties by virtue of which o^*_{ki} refers to some part a_i of object a. (2) Each simple or complex feature (predicate-content) β_j abstractively embodied in o^*_{ki}'s activation state signifies some property B_j such that o^*_{ki}'s-having-β_j represents a_i as having B_j. (3) If o^*_{ki}'s-having-activation-feature-β_j represents a_i as having property B_j, and β_j in turn embodies a higher-level abstraction β'_j, then there is some property B'_j, abstractively embodied in B_j and signified by β'_j, such that o^*_{ki}'s-having-β'_j represents a_i as having B'_j. (4) More generally, let $\{o^*_{ki} : i \in i_m\}$ be a partition of o^*_k into subregions while $\{\beta_i : i \in i_m\}$ is a corresponding array of activation states or features thereof such that, given the structural relations among o^*_k's subregions $\{o^*_{ki}\}$, the array of events $\{o^*_{ki}$'s-having-$\beta_i : i \in i_m\}$ abstractively entails that o^*_k has an activation-state feature β which signifies property B. Then if $\{B_i : i \in i_m\}$ comprises the properties respectively signified by $\{\beta_i\}$, the relevant

structural relations among $\{\underline{o}^*_{ki}\}$ correspond to structural relations among $\underline{a}$'s parts $\{\underline{a}_i\}$ respectively referred to by $\{\underline{o}^*_{ki}\}$ given which the collection of events $\{\underline{a}_i$'s-having-$\underline{B}_i : \underline{i} \in \underline{i}\}$ abstractively contains $\underline{a}$'s-having-$\underline{B}$.

[[Notes: PDD Clause 4 says simply that $\underline{o}^*_k$'s-having-β depicts $\underline{a}$'s-having-$\underline{B}$ only if the more basic events $\{\underline{o}^*_{ki}$'s-having-$\beta_i\}$ constituting $\underline{o}^*_k$'s-having-β respectively depict events $\{\underline{a}_i$'s-having-$\underline{B}_i\}$ from which $\underline{a}$'s-having-B is constituted. But its wording studiously evades details of how structural relations among the chosen parts of $\underline{o}^*_k$ on one hand, and those of $\underline{a}$ on the other, figure in this story. Clause 3 is the special case of Clause 4 wherein $\{\underline{o}^*_{ki}\}$ comprises just $\underline{o}^*_k$ itself.]]

This *principle is tagged with a truth-suspension asterisk because not merely is it contentious whether propositional representation is ever depictive at all, neither is there evident reason why a sophisticated theory of depiction cannot put qualifications on Clauses 1-4 in light of deepened insight into the nature of nominal reference and predicate signification. In particular, inner-pictures must surely be granted a "grain" threshold such that parts of depiction site $\underline{o}^*_k$ which are sub-grain in size are exempted from PDD requirements. Until such time as we discover what qualifications are appropriate, however, we can take PDD as given to idealize intuitive prerequisites for a manner of representation to count as "depiction."[50]

---

[50] PDD is deducible from defining depiction in terms of isomorphism between micro-events sufficient to constitute $\underline{a}$'s-having-$\underline{B}$ and micro-events constituting $\underline{o}^*_k$'s-having-β. But there may also be less extreme versions of inner-picturing, not so simply definable, that also merit consideration as models of depictive representation.

Meanwhile, it should be clear why representations governed by PDD are dense. For if $\underline{a}$'s-having-$\underline{B}$ is depictively represented in $\underline{o}$'s perceiving by $\underline{o}^*_k$'s-having-β, then a rather large proportion of the perceivable factive concommitants of $\underline{a}$'s-having-$\underline{B}$ should be represented simultaneously by $\underline{o}$'s percepts. Moreover, all these depictions take place within the same brain-region $\underline{o}^*_k$, one nested in or more generally constituted out of others, in principle leaving many other regions $\{\underline{o}^*_j\}$ in the remainder of $\underline{o}$'s perceptive brain-stage free to be sites of other inner-pictures in $\underline{o}$.

In contrast, the inner-sentence model tells a very different story of perceptual density and segregation. For if a's-having-B is represented by o's brain-region pair $\langle o^*_{a1}, o^*_{b1} \rangle$ being structured as a subject/predicate frame with the activation states of o^*_{a1} and o^*_{b1} respectively embodying representations α of a and β of B, then o generally needs a different brain-region pair $\langle o^*_{a2}, o^*_{b2} \rangle$ to contain representation of a_i's-having-B_j even when a_i is part of a. (Exceptions will be acknowledged in a moment.) So the inner-sentence model is representationally sparce in holding that when o's perceiving that-α-is-βish represents a's-having-B, o's simultaneous that-α_i-is-$\beta_j$ish perceiving of any factive concommitant a_i's-having-B_j of a's-having-B, if present at all must generally arise from a causal process which at some step splits off from and becomes parallel to the sequence productive of o's perceiving that-α-is-βish. Even so, it is not mandatory for an inner-sentence model of o's simultaneous seeing that-α-is-βish-and-that-α_i-is-$\beta_j$ish to separate these percepts completely no matter how closely connected are the objects a and a_i referenced by nominal concepts α and α_i, or how overlapping are the properties B and B_j signified by predicate concepts β and β_j. In particular, if the α-concept and α_i-concept, or similarly β and β_j, are structurally complex with components in common--e.g., if α is John-and-Mary while α_i is John, or β is rectangularly-blue while β_j is blue--the brain regions that respectively contain α and α_i, or β and β_j, are allowed to have subregions in common containing the shared content. And the counterpart of PDD Clause 3 is arguably true of inner sentences as well. That is, some versions of this model may admit a limited density of perceptual representation in that an inner-sentence which represents object a as having property B perhaps abstracts into arbitrarily many representations of a as having various higher-level properties supervenient upon B. Indeed, inner-sentence theory needs something like that if, as we provisionally accepted earlier though are still prepared to retract, seeing that-this-is-dark-blue entails seeing-that-this-is-blue.

[[For reasons sketched later (p. 278), however, any constitutional model of perceiving generally does best to put sharp constraints on the abstractive nesting of contents it allows. A prudent inner-sentence model would strive for such constraint along lines something like the following: When the inner-sentence embodiment of o's-seeing-that-p consists of o's disjoint brain-region tuple $\langle o_a^*, o_b^*, \ldots \rangle$ being structured as a propositional frame over which that-p's concept elements are distributed as distinctive activity patterns in the various o_i^* (i = a,b,...), only one abstraction γ_i from each o_i^*'s activation state satisfies the criterion (whatever that may be) for γ_i to be a concept. Then an abstraction from the macro-event of $\langle o_a^*, o_b^*, \ldots \rangle$'s total structural/activational condition is an inner-sentence perceiving nested in o's-seeing-that-p just in case it is $F'(O^*) \& \Gamma(O^*)$ for some subtuple $\Gamma(O^*)$ of the events $\langle \gamma_a(o_a^*), \gamma_b(o_b^*), \ldots \rangle$ and a complex F' of structural properties whose possession by subtuple O^* of $\langle o_a^*, o_b^*, \ldots \rangle$ qualifies O^* as a propositional frame in its own right. For example, suppose that o's-seeing-that-this-apple-is-brown-and-bruised has inner-sentence embodiment $F(o_a^*, o_{b1}^*) \& F(o_a^*, o_{b2}^*) \& \alpha(o_a) \& \beta_1(o_{b1}^*) \& \beta_2(o_{b2}^*)$ wherein α, β_1, and β_2 are the concepts this-apple, brown, and bruised, respectively, and $F(\_,\_)$ is the structure of a monadic-predication frame. Then this perceiving also contains perceivings $F(o_a^*, o_{b1}^*) \& \alpha(o_a^*) \& \beta_1(o_{b1}^*)$ (i.e., o's seeing that-this-apple-is-brown) and $F(o_a^*, o_{b2}^*) \& \alpha(o_a^*) \& \beta_2(o_{b2}^*)$ (i.e., o's seeing that-this-apple-is-bruised), but no others. As for our running adjectival test case, whether seeing-that-this-book-is-dark-blue contains seeing-that-this-book-is-blue under our posited constraint on inner-sentence nesting depends on whether this compound predicate embodies dark in a site disjoint from that of blue. That might be so for some predicate modifiers verbalized by adjectives but not for others.]]

Even if inner-sentence predicates do provide densely nested representations of external abstraction hierarchies, however, it is still possible that inner-

sentences can perceptually segregate different levels of abstraction. To bring out this point's essence with minimal distraction, let us momentarily assume (much more strongly than necessary) that if the activation state of any predicatively structured brain-region o_b^* properly abstracts into a pattern β signifying an external property B, then B abstracts from some more determinate property B* signified by o_b^*'s total activation state. Then inner-sentence perceiving is able to segregate levels of property-abstraction if in general, when the total activation state β^* of o's brain region o_b^* signifies an external property B* while an abstraction β from β^* signifies a higher-level property B embodied in B*, o's brain is also able to contain a predicatively structured region o_{bj}^*, distinct from o_b^*, whose total activation state signifies B without signifying any more determinate external property. Commonsense intuitions about the "information loss" that progresses as pre-perceptual input evokes percepts which in turn produce more central cogitations and occasionally eventuate in verbal reports make clear that any reputable model of thinking must allow higher-level predicate concepts to be detached somehow from lower-level embodiments thereof. But whether abstraction levels can be segregated only by successive stages of post-perceptual ideation or whether these can instead occur in perceptual parallel, either simultaneously or as competetive alternatives akin to my potential uttering either 'This is dark blue' or 'This is blue' but not both, is a psychonomic issue that remains widely open.

Flexibility of representation.

We have already noted that inner-picture models of perception draw a hard line between, on the one hand, concepts that can occur predicatively in a percept's content and, on the other, whatever is characterized by the subject-phrases in ordinary-language descriptions of propositions. In prospect, therefore--though how that potential is realized depends greatly on details of our still-to-come theory of aboutness--inner-sentences should be capable of representing many external facts (though not of course all at once) beyond the representational reach of inner-picturing.

Thus when a is a rectangular blue book, if an inner picture can represent shape only nominally by some structural feature of the picture's locus, $\underline{o}$*'s-having-β can depict a's-being-blue by embodying the proposition that-this-rectangular-thing-is-blue but, unlike an inner sentence, cannot depict a's-being-rectangular by embodying a proposition in which the rectangular-concept occurs in the percept's predicate. And neither can perceiving that-this$_1$-yellow-is-more-intense-than-this$_2$-green be a depiction when its this$_1$-yellow and this$_2$-green nominals refer not to physical objects but to particular shades of color displayed nearby. For representation of $\$_1$ish-yellow and $\$_2$ish-green by the activation states of $\underline{o}$'s brain regions $\underline{o}^*_k$ and $\underline{o}^*_j$ would give $\underline{o}$ an inner picture only of one object's being $\$_1$ish yellow joined by another's being $\$_2$ish green.

Moreover, strong limitations on the range of predications available to inner-picturing are also imposed by Clause 4 of PDD. For what that says is essentially that any molar property signified by the global activation state of a depiction locus $\underline{o}^*_k$ must be constituted out of whatever properties are variously signified by the local activations in $\underline{o}^*_k$'s subregions. (If overt speech were like this, we could not truthfully assert 'Those are cattle' unless we refer to something that is partly feline.) In contrast, when $\underline{o}^*_b$ is the predicate locus of an inner-sentence frame, even though $\underline{o}^*_b$'s activation state is constituted by the assorted activities in $\underline{o}^*_b$'s parts, these subregion states are not themselves required to represent anything (though some may do so if $\underline{o}^*_b$'s relevant structural properties make it syntactically complex) and hence place no inherent constraints on what the activation state of $\underline{o}^*_b$ as a whole can signify.

Demonstratives and the targeting of perceptual nominals.

Reluctantly but resolutely, I must advise you to pass over this subsection (pp. 249-274) unless your interest in the semantics of demonstratives or the logic of depictive representation is much deeper than casual. The issue of demonstratives is a large digression from this book's objectives that I would

gladly forego. But it obtrudes in almost every real-life verbal report of propositional perceiving and raises disturbing questions about the extent to which perceptual information can be communicated even in principle, especially by inner-pictures. Like a toad in the soup kettle, once noticed it cannot be ignored but must be disposed of.

(There is, to be sure, a classic way to avoid detecting the demonstrative toad in the first place. That is to take as our paradigm for percept-description the mixed objective/intentional form '$\underline{o}$ sees $\underline{a}$ as βish', wherein '$\underline{a}$' refers not to any component of $\underline{o}$'s mentation but to an object that $\underline{o}$'s percept is about. Yet paraphrasing this form as '$\underline{o}$ sees of $\underline{a}$ that it is βish' makes clear that percepts so described still require $\underline{o}$'s perceiving to include a propositionally structured thought which perforce contains <u>something</u> to bring off reference to $\underline{a}$. Commonsense likes the '$\underline{o}$ sees $\underline{a}$ as βish' form precisely because it allows us to conceive of $\underline{a}$ in any way we fancy without concern for how $\underline{o}$'s percept does this.)

Unhappily, despite the facile simplifications with which I shall dispach this matter, the account is still so long that if you become caught up in it you will have quite forgotten the main currents of this chapter's development by the time those return. But if you will later allow me to treat the subject-content in $\underline{o}$'s-seeing-that-α-is-βish as paradigmatically having composition $\alpha = \underline{T}_{\alpha} \& \mathcal{K}$ without saying much about its nature beyond that $\mathcal{K}$ is a predicative concept nominalized by some inner-syntax adjunctive $\underline{T}_{\alpha}$ of a "target marker" sort hypothesized to underlie such English locutions as 'this $\mathcal{K}$-thing', 'the $\mathcal{K}$-thing', and 'a $\mathcal{K}$-thing', you can skip directly to p. 274 without essential loss of continuity. (Do return to these passages eventually, however. The position they develop is rather important for the theory of representation even if you find it objectionable.)

Our two idealized models of percept constitution also clash instructively in their paradigms for the character of whatever is expressed in perceptual reports by demonstratives. This question proves to be a wonderous snarl of multi-tendriled issues whose sorting out much diminishes the initially large apparent divergence between their inner-picture and inner-sentence accounts. But model contrasts will not be our main concern for some time. Eventually we want to appraise the relative merits of inner-pictures vs. inner-sentences for doing the work that folk psychology expects of thoughts. But first we had better put some perspicuity into the representational intricacies underlying our use of demonstratives.

Ordinary-language efforts to report the propositional details of perceivings almost always include demonstratives--'this', 'that', 'those', 'here', 'there', 'now', 'I', 'you', 'it', 'us', 'them', etc.--in their locutions for perceptual contents. There is evidently something special about the semantics of such terms, as shown by the fancy linguistic footwork incurred whenever we attempt to share knowledge of perceivings so described. For example, suppose that I want to tell you about an overly observant subject in a deception experiment I have been running. If I state

(60-1) **John said, 'I see that your apparatus is misrecording my score'** ,

reproducing therein the sentence that John himself used to convey his observation, I am telling you only what words John uttered, not what I infer from these to have been his percept. To describe the latter, I need something like

(60-2) John saw that my apparatus was misrecording his score ,

whose that-clause differs from John's own percept-description both in its tense (temporal demonstrative) and its reversal of the personal pronouns. And for you to assimilate this information, you must in turn recast (60-2) into, say,

(60-3) John saw that Rozeboom's apparatus was misrecording John's score ,

which eliminates pronouns in favor of a that-clause that preserves representation of (more or less) the same deception event witnessed by John only at the price of considerable departure from the subject-content in John's own perception thereof. Note in particular that (60-2) and (60-3) sacrifice reference to the specific moment in time, verbalized by the tense of John's report, that was now for his perceiving. Were we

to try for that same temporal specificity, we would need to expand (60-3)'s that-clause into something like 'Rozeboom's apparatus misrecorded John's score at 2:17 p.m., May 7th, 1985', exploiting therein chronometric concepts which played no role in John's own awareness then.

The hallmark of demonstratives in perceptual description, then, is that they travel poorly if at all. But is this a genuine cognitive phenomenon or merely a practicality of surface language comparable to the ambiguities and context dependencies so often found with other words? If it merely illustrates the "anaphoric" use of demonstrative terms as local synonyms for non-demonstrative phrases uttered elsewhere, it would have little to do with the nature of perception. And although perceptual reports can seldom be freed of demonstratives by paraphrase in any public language, that may only show the expressive poverty of extant social communication systems. So a useful foil for debate in this matter is the contention that in principle, were we to develop an ideal language containing an unambiguous word or phrase for every concept we are capable of thinking, we could describe all our perceivings without resort to demonstratives.

Now clearly this eliminability thesis has some merit. When I see that-this$_1$-yellow-is-more-intense-than-this$_2$-green, for example, only my meager color vocabulary precludes my verbalizing this by a context-free content clause of form 'that-\$$_1$ish-yellow-is-more-intense-than-\$$_2$ish-green' with as much precision as my self-reports ever manage. Each nominal component of this proposition appears to be a concept that can be repeated in arbitrarily many different perceivings while referring throughout to the same determinate shade of color. From there, it is straightforward to envision models according to which your hearing me utter 'I see this yellow as more intense than this blue', together with your observing my gestures and the colors of nearby objects, evokes in you the opinion that-it-appears-to-this-guy-that-\$$_1$ish-yellow-is-more-intense-than-\$$_2$ish-green, or its kin, wherein the \$$_1$ish-yellow and \$$_2$ish-green concepts are copied (nevermind how accurately)

out of your current perceptual experience into your judgment about how these colors appear to me. Your $\$_1$ish-yellow and $\$_2$ish-green concepts so recruited may well differ appreciably from the corresponding components of my own percept, but they can also be nearly the same if our perceptual mechanisms are similarly tuned.

More commonly, however, paraphrastic eliminability of demonstratives from perceptual reports seems dubious even with the resources of an ideal language. When John has the percept he reports as quoted in (60-1), its component representations of the Rozeboom-stage and John-stage that he expresses by 'your' and 'my' may indeed include conceptual ingredients describable by English adjectives. (E.g., balding, messy-labcoat, and peeking-surreptitiously-over-clipboard might be fragments of John's momentary you-concept.) Yet beyond that, an essential facet of these percept components is prima facie simply their being there, rather than somewhere else, and accomplishing reference thereby in some fashion fundamentally different from that of concepts whose referents are retained across repetitions. Most starkly this seems true of the now-concept expressed by the tense of John's report; and the same nearly featureless indexicality--a bare "deictic" function (Lyons, 1977, p. 637f.)--can be seen in other percepts that represent spatial locations by contents verbalized as 'here' and sometimes an unqualified 'this'. First-person pronouns also intuit as largely deictic, albeit what these pick out on their various occasions of usage remains enigmatic.

The mystery of demonstratives begins with their operation as linguistic devices and cannot be fathomed apart from some psychonomic account of verbal communication. In particular, whatever is context-sensitively distinctive about how demonstratives convey ideation needs to be separated from whatever may be special in what they convey. The latter, our main concern here, turns on the extent to which mental representations can be shared; and I shall sketch only such idealized fragments of the communicative process as are required to anchor that pivot.

The first of these fragments posits that when I verbalize a sentence 'o_1 sees that p' under standard communicative circumstances, my aim is (a) to inform some hearer o_2, say you, about a certain perceptual representation activated in observer o_1, say me, and (b) to do so moreover by evoking in you, as part of the message conveyed, a simulacrum of the representation in me (i.e. o_1) that this message is about. Let o_1^*'s-having-P_1 be the t-core of my to-be-communicated seeing-that-p, where o_1^* is my brain site for the propositionally patterned property P_1 that constitutes my perceiving's content. (More technically, take P_1 to be the t-core pattern in the thinking-that-p which abstracts from my more modally determinate seeing-that-p.) And let o_2^*'s-having-P_2 be that part of your message-induced awareness of my-seeing-that-p in which your simulacrum of my perceptual representation is localized. Then the more closely o_2^*'s-having-P_2 resembles o_1^*'s-having-P_1 in all representationally relevant respects, the more ideally I have communicated my seeing-that-p to you.

[Saying in this case that I aim to give you a "simulacrum" of my that-p thought condenses two important points that a serious study of communication would probe in detail. First, when I utter a sentence 'q' to you, it is almost always my intent (latent if not phenomenally conscious) to activate in you a thought whose representational character is similar, in major albeit still obscure respects, to my own active ideation for which 'q' is an expressive vehicle in my language. (This is true even when, deceitfully, I endeavor to give you a mode for this shared thought-content different from mine.) And secondly, when my utterance has embedded-sentence form 'o_1 sees that p' (where 'sees' can just as well be any other Psi-verb), a prominent component of the propositional content that-p_1^+ this elicits in you is essentially the same proposition that-p_1 you would have received had I uttered just 'p'. (That oversimplifies a bit when 'p' contains demonstratives and 'o_1' is not first-person singular--cf. (60-2)--but it captures the gist of standardly intentional

Psi-verb communication.) This embedded proposition is (part of) what my communication as understood by you is *about*; but your received information *that*-p_1^+ contrives to designate the *that*-p_1 idea by literally *exhibiting* it within your representation of me as seeing-that-p_1. This is a remarkably special style of reference that is possible only when the object represented is mental, and even then is not available for communication about most thoughts which a technical science of mind might seek to study. Thus in particular, were you and I privy to a specialist vocabulary $\{'P_k'\}$ of percept identifiers defined by advanced perceptual theory in terms of their purported referents' distinctive psychonomic functions, my telling you 'o_1 has percept P_k' should evoke in you ideation that simulates the information state in me that initiates this communication (Point One), but does *not* thereby induce in your thinking either a literal recurrence of P_k or any functional near-equivalent thereto (contra Point Two). Only in the semantically atypical but commonsensically prominent case of communicating mental states by *that*-clause completions of Psi-verbs does your received information ideally include a simulacrum of the representation represented.[51]]]

---

[51] *A prospectus in the philosophy of mind*. Suitable expansion of this point, that a technical science of mentality seeks to talk about thought contents *without* in general representing them by simulacra such as invoked by our commonsense language of intentionality, largely resolves various problems of "subjective qualia" which are so often alleged to defeat functionalist/materialist accounts of mentality. But here is not the place to develop that claim.

But what are these "representationally relevant respects" which make for ideal sharing of thoughts and motivate describing your received information about my perceptual representation as containing a "simulacrum" of the latter? Above all, one is my percept's conceptual content while another is its factive object. And what makes information sharing so tricky is that our agreeing in one of these respects may preclude agreement in the other.

Specifically, two communicative ideals are possible here, the real-world incidence of near-approaches to which need not concern us. Giving name to the first, let us say that your o_2^*'s-having-P_2 simulacrum of my o_1^*'s-having-P_1 representation is (ideally) sympathetic iff $P_2 = P_1$. That is, stated loosely, your reproduction of my perceiving is sympathetic iff its ideational content P_2 is the same as my perceptual content P_1. (Do not protest that your P_2-thought could never approach the sensuous quality of my P_1-percept. Commonsense insists that perceiving-that-p, hoping-that-p, surmising-that-p, etc., have something that-p-ish in common; and for present purposes it does no harm to presume that even if the complete activation state of my perceptive brain-region o_1^* cannot be reproduced in your o_2^*, P_1 is an abstraction from o_1^*'s total state that also abstracts from certain states possible for o_2^*.) And for the other communicative ideal, say that your o_2^*'s-having-P_2 simulates my o_1^*'s-having-P_1 objectively iff what these two mental events are respectively about is the same for each. That is, when the factive referent of my seeing-that-p is a's-having-B, the simulacrum of this evoked in you by my utterance 'I see that p' is objective if it too represents a's-having-B. For you to be perfectly informed about my perceiving, we would like your simulacrum to be both sympathetic and objective. But to what extent is that possible?

The answer turns on whether o_1^*'s-having-P_1 accomplishes representation solely through the conceptual pattern therein, or whether the entire event is required. If it is just the propositionally structured property P_1 which represents a's-having-B (or represents a as having B), regardless of where that pattern occurs, then if P_2 in your o_2^*'s-having-P_2 simulacrum of my o_1^*'s-having-P_1 is identical with P_1 your evoked thought pattern P_2 (= P_1) is not merely a sympathetic repetition of my percept's content but also represents in you the very same factive object, a's-having-B, perceptually represented in me.[52] However, an alternative prospect

[52] When we here posit that a thought-pattern μ (propositional, predicative, or nominal) stands for an objective entity e (event, property, or particular), we shall for simplicity speak as though μ's representation of e in any particular o_k^*'s-having-μ

instantiation thereof is strictly a binary relation of μ-ness to $\underline{e}$ for which additional features of $\underline{o}_k^*$ and its surround are irrelevant. But that is an enormous idealization; for whatever may be the nature of any aboutness coupling between μ and $\underline{e}$, it must surely reside to some extent in dispositional properties of the system containing $\underline{o}_k^*$ by virtue of which μ functions as it does in this system. That is, μ stands for $\underline{e}$ not *simpliciter* but only relative to certain domain-stable support conditions $\underline{C}_{\mu,e}$ that can be viewed as part of the domain preconditions defining some semantically specialized *kind* of intentional system. (Indeed, rather than say that μ and other thoughts are "meanings," as is my wont, it can be argued instead that "meaning" is something that thought μ *has* in system $\underline{o}$, namely, the functional role in $\underline{o}$ characterized by $\underline{C}_{\mu,e}$.) Even so, representation of $\underline{e}$ by μ under $\underline{C}_{\mu,e}$ is still a *patternwise* aboutness that can recur repeatedly in systems of this kind, and which can be described as a binary relation by saying that what represents $\underline{e}$ is not just μ but the more global pattern property, μ-activated-in-the-context-of-domain-constraints-$\underline{C}_{\mu,e}$. But you don't want to be burdened with repeated mention of such complications, especially since we shall have nothing useful to say about them.

---

is that what represents $\underline{a}$'s-having-$\underline{B}$ in my perceiving is not just propositional

content $\underline{P}_1$ qua repeatable pattern, but the full $\underline{o}_1^*$'s-having-$\underline{P}_1$ perceptual *event* in such fashion that when this same $\underline{P}_1$ recurs in another location $\underline{o}_j^*$, the factive referent (if any) of $\underline{o}_j^*$'s-having-$\underline{P}_1$ is generally some event other than $\underline{a}$'s-having-$\underline{B}$. If this second alternative is how representation works, then the only way for your $\underline{o}_2^*$'s-having-$\underline{P}_2$ to be an objective simulacrum of my $\underline{o}_1^*$'s-having-$\underline{P}_1$ is for your thought-content $\underline{P}_2$ to differ from my $\underline{P}_1$ in some fashion that manages to give your thought-*event* $\underline{o}_2^*$'s-having-$\underline{P}_2$ the same factive referent as my perceptual event $\underline{o}_1^*$'s-having-$\underline{P}_1$ by compensating for $\underline{o}_2^*$'s displacement from $\underline{o}_1^*$.

Why a complete mental event, not just the activation state therein, may be required for factive representation is plain in the inner-picture model of perception. For as already noted, this envisions that an internal depiction of $\underline{a}$'s-having-$\underline{B}$ is some brain event $\underline{o}_k^*$'s-having-β wherein, even though activity pattern β putatively signifies $\underline{B}$-ness regardless of where β may occur, the particular object $\underline{a}$ here represented as having $\underline{B}$ is designated by depiction site $\underline{o}_k^*$ on the basis of this situation's locus structure. Presumably, the facet of structure most salient for selecting $\underline{o}_k^*$'s referent is $\underline{o}_k^*$'s position in space-time. (Or at least we can let that go proxy for some more complicated story about $\underline{o}_k^*$'s functional positioning in a neural network.) So one simple example of how $\underline{o}_k^*$ might pick out an object $\underline{a}$

to be represented as having B by o_k^*'s-having-β is for this a to be whatever region of space-time is the shape, size, and distance from o_k^* identified by a certain fixed function of o_k^*'s own shape, size, and orientation to the major axes of the observer-stage whose brain contains o_k^*. (Nevermind how silly this particular schema for o_k^*-to-a reference may seem; it merely illustrates how it is possible for o_k^*'s physical geometry to select an external target of representation, not what may be a plausible depictive account of this.) Accordingly, if o_k^*'s-having-β so represents a as having B, and o_j^* is some other brain region (notably, in some observer-stage other than the one containing o_k^*) that is structurally just like o_k^* except for location, then the object represented as having B by o_j^*'s-having-β would be not a but some other thing geometrically related to o_j^* in the same way that a is related to o_k^*. In special cases, it may be possible for o_j^* too to refer to a if the shift in position from o_k^* to o_j^* is suitably compensated for by o_j^*'s also differing from o_k^* in other structural respects; but in general it should be difficult if not impossible for any o_j^* widely separated from o_k^* to refer depictively to the very same a represented by o_k^*.

In short, as portrayed so far, depicted information is virtually incommunicable. For reproduction of an inner-picture's pattern in different locations should generally fail to preserve factive reference, and it is dubious how often common reference can be achieved by varying pattern across different depictive events.

[[In the variant of inner-picturing that holds the representation of external object a in o_k^*'s-having-β to be not o_k^* itself but the structural condition F_α of o_k^* by virtue of which the first version of depiction takes o_k^* to designate a, we can say that what represents a as having B is not strictly the event o_k^*'s-having-β but only the complex property, F_α-coexemplified-with-β. But the latter is not a repeatable pattern which can be communicated. For so long as F_α includes the property of having location thus-and-so, F_α and hence F_α's conjunction with any β can occur only at site o_k^*. When the theory of aboutness contrasts representation-

by-events with representation-by-patterns, it is only repeatable patterns capable of participation in lawful regularities that we accept for the latter.]]

In contrast, inner-sentence representations are paradigmatically repeatable patterns even though it is not precluded that some may be uncommunicable events. Consider the simplest form of primary perceiving, an $\underline{o}$'s-seeing-that-α-is-βish wherein truthful representation of some nonrelational $\underline{a}$'s-having-$\underline{B}$ event is contained. In its inner-sentence construal, the t-core of this perception consists of a pair $\langle \underline{o}_a^*, \underline{o}_b^* \rangle$ of $\underline{o}$'s brain regions satisfying whatever complex $\underline{F}$ of repeatable structural conditions establishes $\underline{o}_a^*$ and $\underline{o}_b^*$ as respectively the subject-position and predicate-position of a monadic propositional frame, while certain abstractions α and β from the activation states of $\underline{o}_a^*$ and $\underline{o}_b^*$ are the α-concept and β-concept, respectively. (We needn't make explicit here that even for a nonrelational β, $\underline{o}_a^*$'s-having-α may consist in various subregions of $\underline{o}_a^*$ having a certain configuration of structural relations and component activations, reflecting the grammatical complexity of an ordinary-language verbalization of the α-concept.) Then the propositional content of this perceiving is the pattern-property $\underline{P}_{\alpha\beta}$ on brain-region pairs such that, by definition, any $\langle \underline{x}, \underline{y} \rangle$ has $\underline{P}_{\alpha\beta}$ iff $\underline{F}(\underline{x},\underline{y}) \,\&\, \alpha(\underline{x}) \,\&\, \beta(\underline{y})$. Pending deeper insight into the nature of predicate semantics, we continue to presume that external $\underline{B}$-ness is (or can be) signified by internal activation pattern β wherever this may occur (albeit recall fn. 50, p. 254). So if we can top off the inner-sentence model of representation with an account of nominal reference under which α qua repeatable pattern picks out $\underline{a}$ as referent, $\underline{P}_{\alpha\beta}$ represents $\underline{a}$'s-having-$\underline{B}$ wherever this proposition is instantiated, not only by $\langle \underline{o}_a^*, \underline{o}_b^* \rangle$ but also by any other pair $\langle \underline{o}_i^*, \underline{o}_j^* \rangle$ of brain regions for which it holds that $\underline{F}(\underline{o}_i^*, \underline{o}_j^*) \,\&\, \alpha(\underline{o}_i^*) \,\&\, \beta(\underline{o}_j^*)$.

How the α component of propositional pattern $\underline{P}_{\alpha\beta}$ might designate a particular localized object $\underline{a}$ is plain in the commonsense semantics of definite descriptions. Consider a nominal of form 'the $\underline{Q}$'--e.g., 'the largest city in Europe', 'the first moon landing', 'the s.o.b. who stole my raincoat last week', etc.--wherein '$\underline{Q}$' is

a more-or-less complex predicate which, contrary to the third example just given, we shall presume to be free of demonstratives. (In practice, definite descriptions often suppress article 'the' in favor of a possessive construction, as in 'Europe's largest city', 'Brahms' 4th symphony', etc.) Commonsensically, this locution refers to whatever object has the property signified by '$\underline{Q}$' so long as there is exactly one such thing; and that is our first choice for idealizing how inner-sentences achieve nominal reference. Specifically, we posit as a first-approximation that the nominal component α of inner-sentence propositional pattern $\underline{P}_{\alpha\beta}$ is a compound $\alpha = \underline{T}_{\alpha} \& \chi$ (i.e., $\alpha(\underline{x}) = \underline{T}_{\alpha}(\underline{x}) \& \chi(\underline{x})$) wherein an activity pattern χ signifying some generally-complex property $\underline{K}$ is co-exemplified with a special repeatable "target-marker" subpattern $\underline{T}_{\alpha}$ whose function includes (inter alia) the inner-syntax role externalized by the word 'thing' in 'thing that is $\underline{K}$-ish'. (Without $\underline{T}_{\alpha}$ or some other target-marker to augment χ in α we should prima facie view α as referring not to some particular object of kind $\underline{K}$ but to the property $\underline{K}$ itself--whence $\underline{P}_{\alpha\beta}$ would become a representation of $\underline{K}$-ness as having $\underline{B}$. We shall say more about target-markers shortly.) Then in a classical theory of reference (which however needn't be the last word on this matter), α refers to object $\underline{a}$ iff $\underline{a}$, and $\underline{a}$ alone, has property $\underline{K}$. And if α does so refer just to $\underline{a}$, then wherever propositional pattern $\underline{P}_{\alpha\beta}$ occurs it represents $\underline{a}$ as having the property $\underline{B}$ signified by $\underline{P}_{\alpha\beta}$'s predicative component β.

Large semantic-theoretical questions bedevil the model of nominal reference just sketched (which is not to say that any other psychonomically honest approach to this has it any easier), especially regarding the representational status of α (= $\underline{T}_{\alpha} \& \chi$) when many things have the property $\underline{K}$ signified by χ. To be sure, we can hope that this circumstance seldom arises for inner-sentence percepts. For if primary perceptual contents are far more richly determinate than perceptual reports ever verbalize except by demonstrative allusion to attributive concepts newly minted on each perceptual occasion, and if, moreover, arbitrarily much

of this detail can be packed into the percept's subject-content α (as the inner-sentence model is free to presume), then it could well result that the attributive subpattern χ which delimits α's referential target signifies a property so elaborately determinate that scarcely ever would it be exemplified by more than one real-world object. Even so, it seems foolhardy to make so strong a uniqueness presumption foundational in our theories of perceptual aboutness. In particular, even if each demonstrative construction of form 'this K-thing' manages to go linguistic proxy for a different concept on each occasion of its usage in perceptual reports, introspection finds it implausible that this is also true of 'I', 'here', and 'now'.

So what might α represent if there are several K-kind objects in the world? One radical prospect is that α multiply designates each object a of kind K, while $P_{\alpha\beta}$ then represents a as having the property B signified by the β-component of $P_{\alpha\beta}$ just in case α designates a and a has B. According to this proposal, $P_{\alpha\beta}$ is veridical just in case some K-thing has B, yet what $P_{\alpha\beta}$ represents is each determinate a's-having-B event wherein a is a K that has B. (Any such multiple-representation thesis is a major break with orthodoxy in philosophical semantics. Nevertheless, I have elsewhere argued repeatedly--e.g., Rozeboom, 1970b; see also Rozeboom, 1973--that to comprehend semantic reality it is essential that we abandon idealizing reference as a word-to-world mapping under which nominals refer uniquely or not at all.) Alternatively, we can waive reference for α qua pattern, and hold instead that when the target-marker T_α which modifies χ in $\alpha = T_\alpha \& \chi$ includes one or more marker components of a special "token-cue" sort (to be elaborated shortly), the event of α's being instantiated in some particular brain region o_a^* designates whatever K-kind object is causally closest to o_a^* (or is most salient for o_a^* in some other respect selected by the particular token-cues in T_α.) Token-cues are repeatable pattern components having a bare deictic function, as verbalized most purely by hard-core demonstratives 'I', 'here', and 'now'. Under this latter treatment, whether an inner sentence

accomplishes representation patternwise by its t-core content $\underline{P}_{\alpha\beta}$, or only eventwise by particular instantiations of $\underline{P}_{\alpha\beta}$, depends on whether $\underline{P}_{\alpha\beta}$'s subject-content includes token-cues. If it does not, then and only then does $\underline{P}_{\alpha\beta}$ aspire to represent some a's-having-B event qua repeatable pattern and thus to constitute communicable information--which, however, still leaves the problem of what such a $\underline{P}_{\alpha\beta}$ might represent when its subject-content is insufficiently replete to single out a unique referent. We shall make one last pass at this issue following a small shift in perspective.

Demonstrative reference reconsidered: Existential representation.

The possibility of communication by inner-pictures is not quite so bleak as just made out. For arguably, I have been arrogating a canonical form for perceptual representation that is biased against depiction. Inner-pictures can, in principle, transmit existential generalities; and although I have made considerable show of positing primary percepts to have singular subject/predicate form, as distinct from what modern logic takes to be the form of existentially quantified propositions, it is time to acknowledge that ordinary-language usage of indefinite descriptions appears to achieve the force of the latter with the syntax of the former.

To appreciate the subtleties here, let us articulate some structure within the subject-content of seeing-that-α-is-βish by letting χ be some attributive concept which a perfected English could express by an adjectival phrase, and compare

(61-1) o sees that this χ-thing is βish ,

(61-2) o sees that a χ-thing is βish ,

(61-3) o sees that something which is χ is also βish ,

(61-4) o sees that there is something which is both χ and βish.[53]

---

[53]Nothing devious is intended here by taking 'χ' to be adjectival as given while 'β' is adjectivized by a '-ish' suffix. Feel free to treat 'χ' as interchangeable with 'χish', and 'βish' with 'β'.

The that-clause in (61-1) is paradigmatically a singular subject/predicate proposition, whereas in (61-4) it is an existential quantification whose syntax is made clear by rewriting (61-4) in symbolic-logic notation as

(61-4a) o sees that $(\exists \underline{x})[\chi(\underline{x}) \& \beta(\underline{x})]$.

But what about (61-2) and (61-3)? The latter--which I include here merely to illustrate that ordinary language offers more than one grade of intermediary between (61-1) and (61-4)--seems similar enough to (61-4) that we can tolerate its assimilation to (61-4a) despite qualms whether 'Something which is χ is also βish' differs from 'Something which is βish is also χ' by no more than the trivial permutation of coordinate predicatives in '$(\exists \underline{x})[\chi(\underline{x}) \& \beta(\underline{x})]$' vs. '$(\exists \underline{x})[\beta(\underline{x}) \& \chi(\underline{x})]$'. But linguistic intuition cries out against paraphrasing (61-2) as (61-4a); for the latter fails to capture the former's subject/predicate asymmetry. The first predicative in 'A χ-thing is βish' has a manifestly different syntactic role than does the second; and although the psychonomic import of that difference still remains for clarification, we have every reason to anticipate that thought patterns that-a-χ-thing-is-βish and that-a-βish-thing-is-χ have appreciably different internal causes and effects despite their having the same truth-condition, namely, $(\exists \underline{x})[\chi(\underline{x}) \& \beta(\underline{x})]$. The propositional content in (61-2) is as much of singular subject/predicate form as is the content in (61-1), even if our theory of cognitive representation may desire to give one a different sort of factive object than it assigns to the other.

So how do (61-1) and (61-2) differ in their propositional contents? I suggest that there needn't be much difference at all, so long as the attributive content demarked by 'χ' in (61-2) is not limited to what extant English can express. Or, somewhat more broadly, I shall argue that (61-1) abstractively contains (61-2), in the way that seeing-that-this-is-dark-blue may contain seeing-that-this-is-blue, with the two becoming largely the same when the χ-concept is replete with all the attributive content for which the demonstrative goes proxy

in the particular context of (61-1)'s usage. The case for this runs as follows: In any reasonable constitutional model of o's-seeing-that-α-is-βish, the locus o* of this perceiving's t-core, o*'s-having-$P_{\alpha\beta}$, will contain a subregion o^*_a (which is the entirety of o* in the inner-picture model but only a proper part of it in the inner-sentence view) such that the repeatable subject-content in propositional pattern $P_{\alpha\beta}$ is a subpattern α exemplified by o^*_a. This repeatable subpattern α in turn analyzes as a compound $\alpha(x) = T_\alpha(x) \& \chi(x)$ wherein χ is an orthodox though generally complex and only poorly verbalizable attributive concept (i.e., χ is also capable of occurring in the predicate position of a subject/predicate propositional frame) that in principle--not necessarily in fact--signifies some objective property K-ness, while T_α is some subpattern in a special class $\{T_j\}$ of auxillary brain-region features which may be called "target-markers." The latter variously constitute whatever is added to the χ-attributive when that is nominalized by one of the transformations whose most prominent instances are expressed in Enghish by locutions

(62-1) this* χ-thing ,

(62-2) the χ-thing ,

(62-3) a χ-thing ,

(62-4) χ-ness ,

(62-5) χ-kind .

[[Target-marker feature T_α in $\alpha = T_\alpha \& \chi$ might be either structural (and hence optionally assimilable into the syntactic frame of propositions that include it) or activational. And although for simplicity we shall speak as though T_α and χ in o^*_a's-having-α are both properties of o^*_a as a whole, it is alternatively possible and indeed perhaps more likely that the composition of $\alpha(o^*_a)$ is $T_\alpha(o^*_{a_1}) \& \chi(o^*_{a_2})$ for disjoint or at least distinct subregions $o^*_{a_1}$ and $o^*_{a_2}$ of o^*_a. The label "target-marker" alludes to T_α's fine-tuning of the referential aim of nominals in which it occurs.]]

The demonstrative in (62-1) is marked with an asterisk to signal that for present purposes we want this to be read in a purely deictic sense, not as proxy for still more attributive content which could be expressed in a perfected language by additional demonstrative-free adjectival phrases conjoined with χ. That is, for getting at what the demonstrative contributes uneliminably here, we presume that all descriptive content which could in principle be expressed by verbalization of (62-1) on some particular occasion to which this locution's semantic status is relative is already in the χ-concept.

But just what do the differences in array (62) amount to, anyway? Classical semantics answers in terms of what the concepts these phrases express are about, i.e., what they purport to represent under what circumstances: For (62-1), each particular occurrence of this nominal--not the repeatable pattern itself but only some "token" thereof--aspires to designate a particular object of χ-demarked kind K in that occurrence's vicinity. In contrast, the (62-2)-concept purportedly refers qua repeatable pattern, on all occasions of its usage, either to the same one-and-only K-thing or, lacking K-uniqueness, to nothing. As for (62-4,5), each expresses a nominal designed to represent qua pattern a singular universal, namely, a unique property signified by χ in (62-4) and the corresponding class of objects in (62-5). And although classical semantics does not concede nominal reference to the (62-3)-concept, it does proffer objective truth-conditions for propositions containing this construction in subject-position. Yet the more fundamental question, which philosophical semantics traditionally ignores, is what functional distinctions ground these contrasts in representation? That is, how does one nominalization of the χ-predicative differ from another in its arousals' causes and effects?

To argue that the differences among (62-1,2,3) in perception may be minuscule --(62-4,5) are a different story that needn't concern us here--I put it to you that with as much paraphrastic equivalence as real language ever provides, (62-1,2,3) have essentially the same meanings, respectively, as

(63-1) a here-and-now α-thing ,

(63-2a a one-and-only α-thing ,

(63-2b) a uniquely here-and-now α-thing ,

(63-3) a (not necessarily here-and-now) α-thing .

Or rather, any intuitive discrepancy between (62-i) and (63-i) can be written off to ambiguities in the former that are reduced in the latter. Thus, whereas (63-2a) is a lawyer's reading of (62-2), (63-2b) more tightly captures what would ordinarily be understood by (62-2) in a perceptual report and may also be closer than is (63-3) to some occurrences of (62-1). Also, the parenthetical phrase in (63-3) does not voice an active part of the concept standardly expressed by (62-3); rather, it points out what (63-3) must omit if that is to disambiguate (62-3) as distinct from (63-1), albeit this disclaimer may indeed be explicit in some occurrences of (62-3). And we should further note that "here-and-now" is amenable to refinements and modifications which everyday English cannot verbalize precisely but are intimated by such awkward phrasings as 'almost here and now', 'over there a little while ago', etc. (Much of the shading in such demonstrative phrases may well be expressible attributive content; but don't bet that it all is.)

If array (63) is indeed a disambiguating but otherwise accurate paraphrase of array (62), it makes clear that when the concept α_i expressed by (62-i) or (63-i) for each $i = 1,\ldots,5$ is modeled as a repeatable pattern wherein common predicative χ is conjoined with a distinctive target-marker T_i, i.e. $\alpha_i = T_i \& \chi$, these target-markers T_1, T_{2a}, T_{2b}, T_3, as well as other fine variants by which list (63) can be expanded, can all be viewed as the conjunction $T_i = T_0 \& T_i^+$ of a basic marker T_0: a-( )-thing with an enrichment feature T_i^+ which for $i = 3$ is null except when the parenthesis in (63-3) reflects an active disclaimer. That is, I submit not merely that the thought-pattern α_3 expressed by (62-3) is fully as much syntactically nominal as are the undisputed nominalizations α_1, α_2, α_4, α_5 of χ expressed by (62-1,2,4,5), but moreover that α_1 and α_2 are just α_3 with

small enhancements. Yet how commonly perceptual contents occur with basic target-marker $\underline{T}_0$ so enhanced remains problematic. The question is not whether $\underline{T}_0 \& \underline{T}_1^+ \& \chi$ with non-null $\underline{T}_1^+$ can be aroused in a percept's subject-position but how often that is profitable when balanced against costs, and moreover whether some $\underline{T}_1^+$ (notably $\underline{T}_1^+$) shouldn't be treated as features of the modes in which thoughts get entertained rather than as fragments of conceptual content.

Regarding definite-article target-markings, it seems evident that input conditions which evoke seeing-that-a-χ-thing-is-βish would scarcely ever provide epistemic justification for strengthening this into seeing-that-a-one-and-only-χ-thing-is-βish. Starting from (63-1) on the other hand, it could be both advantageous feasible and prospectively to perceive truthfully that-a-unique-here-and-now-χ-thing-is-βish rather than just that-a-here-and-now-χ-thing-is-βish. Yet there is little information gain in this uniqueness addendum unless χ is much attenuated from the nominal content in the most richly determinate seeing-that-α^*-is-βish from which this seeing-that-a-here-and-now-χ-thing-is-βish abstracts, and in that case the uniqueness enhancement becomes epistemically risky unless consequent upon a careful process of perceptual scanning. So by all rights, enrichment of (63-3) into (63-2a) or (63-1) into (63-2b) should in primary perception be an infrequent departure from a uniqueness-noncommital norm. Accordingly, the main challenge of (62/63-1,2,3) lies in what seeing-that-a-here-and-now-χ-thing-is-βish has to gain or lose over seeing merely that-a-χ-thing-is-βish.

What here-and-now loading contributes to a perceiving's functional import is urgency. Arousal of pattern $\underline{T}_0 \& \underline{T}_1^+ \& \chi$ as the subject-content of a percept whose predicate is β is disposed by interaction with the system's memory store to impel immediate effector actions that would be inappropriate in response merely to belief that some χ-thing not here-and-now is βish. Yet that difference scarcely seems to matter in the early phases of input processing; for the sensuous insistence of perceiving that-a-χ-thing-is-βish--i.e., what this particular modality adds

to conceptual contents whose descriptions accept couplings with a plurality of Psi-verbs--is tantamount to a here-and-now content feature; nor in the main, ignoring the gazings of astronomers, is there much chance that o's seeing-that-a-χ-thing-is-βish is veridical unless an object of the $\chi\&\beta$-signified sort is in o's vicinity. So perceptual activation merely of a-χ-thing instead of a-here-and-now-χ-thing should make little difference for what this makes happen next.

In post-perceptual ideation, on the other hand, the presence/absence of T_1^+ should indeed matter--as can be appreciated most easily by appeal to commonsensically idealized memory. Suppose that after a continuant system s passes through some stage o where

(64-1) o sees (veridically) that a here-and-now χ-thing is βish ,

processes of retention and recall initiated by (64-1) bring it about in a later stage o' of s that

(64-2) o' remembers that a $T_\mu^+\&\chi$-thing is βish ,

wherein T_μ^+ either is here-and-now or is null. Clearly, this remembering's import for action in stages of s immediately following o' should be importantly influenced by whether T_μ^+ carries here-and-now urgency. But the appropriateness of that impulsion depends on how far o' is displaced from o. Were (64-2) to be an episode of immediate memory, with o' following o by scant seconds, then T_μ^+ = here-and-now is not merely correct but quite likely requisite for the mediation through which (64-1) prompts the system to take such actions as may be cogent in the presence of an object having the character represented. But if (64-2) is long-term recall, with o' far removed in space/time from o, then its content will be veridical if T_μ^+ is null but is in all likelihood false, with unhelpful if not adverse consequences, if T_μ^+ is here-and-now.

Let us call deictic enrichment feature T_1^+: here-and-now a "token-cue" (cf. p. 258, above) and allow that commonsense variants on this and other hard-core

demonstratives, such as expressed by 'recently nearby' and 'my', suggest that mentation exploits a repertoire of such features, fading in the limit to null or its virtual equivalent somewhere-somewhen. The point to be taken from (64-1,2) is that whatever particular token-cue $\underline{T}^{+}_{\alpha}$ may enhance $\underline{T}_{0}$ as a variant reading of here-and-now in (64-1), it is maladaptive for this very same $\underline{T}^{+}_{\alpha}$ to persist throughout the succession of this perceiving's consequences unless $\underline{T}^{+}_{\alpha}$ is null. And that holds not merely for remembering/recalling this percept's full propositional content as envisioned by (64-2), but also for recombinant retention of its nominal fragment as when, e.g., event (64-1) interacts with $\underline{o}$'s conviction that-most-βish-things-are-γs to evoke conclusion that-a-$\underline{T}^{+}_{\mu}\&\chi$-thing-is-probably-a-$\gamma$. Briefly, when $\underline{T}_{0}\&\underline{T}^{+}_{\alpha}\&\chi$ is the subject-content in a belief (perceptual or otherwise) in $\underline{o}$ that gives rise to a belief with subject-content $\underline{T}_{0}\&\underline{T}^{+}_{\mu}\&\chi$ in successor $\underline{o}'$ of $\underline{o}$, the token-cue trace $\underline{T}^{+}_{\mu}$ of $\underline{T}^{+}_{\alpha}$ that is cognitively optimal for the latter depends on the excursive interval from $\underline{o}$ to $\underline{o}'$, with optimal $\underline{T}^{+}_{\mu}$ fading rapidly to null with increasing lag.

Despite the importance of retaining non-null $\underline{T}^{+}_{\mu}$ or its force-equivalent in the short-term effects of a perceptually activated thought-pattern $\underline{T}_{0}\&\underline{T}^{+}_{\alpha}\&\chi$, I take the point just developed to urge conclusion that token-cues should be excluded from those brain-activity abstractions we classify as conceptual content. Token-cues are by no means to be ignored; but how they work (or ought to work) in rational thinking is so different from the behavior of content components whose character should ideally be preserved under transfer from one thought-event to another that we do best to split off token-cuing from conceptual content and treat this instead as a facet of mentation akin to cogitive mode. Indeed, there is no clear reason why variation in token-cue urgency cannot be straightforwardly subsumed under modal contrasts so long as we recognize that modalities can be nested, one within the scope of another, with their innermost components capable of attachment not just to propositions as wholes but selectively to subpatterns therein.[54]

---

[54]More broadly, there is reason to suspect that many aspects of mentation verbalized in ordinary English by special constructions within the that-clause completions of Psi-verbs are best reconstructed as modal attachments to conceptual content. The modality of thought is a terra incognita far more complexely vast than we have yet dared to acknowledge.

Even so, to sustain our present avoidance of modality issues, let us summarize the psychosemantics of deixis expediently as follows: For any cogitive mode ϕ, whenever it occurs that

(65-1) $\underline{o}$ ϕs that a $\underline{T}^{+}_{\alpha}\&\chi$-thing is βish

for some token-cue $\underline{T}^{+}_{\alpha}$ expressed in everyday English by a deictic demonstrative, it follows by abstraction from (65-1) also that

(65-2) $\underline{o}$ ϕs that a χ-thing is βish .

And conversely, the forcefulness of entertaining concepts in perceptual mode suggests that any instantiation of schema (65-2) with ϕ a variant of perceiving is a-derivative from some instantiation of (65-1) with non-null $\underline{T}^{+}_{\alpha}$. (Whether this converse also holds for modalities other than perceiving is problematic; but the grammatical tenses urged upon their completion clauses by 'remembers' and 'anticipates' point toward a larger story in this regard.) The cogitive consequences of (65-2) for any successor $\underline{o}'$ of $\underline{o}$ are included in those of (65-1), and we may plausibly conjecture that, apart from open-loop feedback, the excess of the latter over the former decreases rapidly to null with increasing lag from $\underline{o}$ to $\underline{o}'$.

[[More technically, we envision that the trace of $\underline{T}^{+}_{\alpha}$ in the effects of (65-1) undergoes short-term changes that are probably rather like a decay to nullity, leaving (65-1)'s long-term cogitive consequences pretty much the same as what they would have been from an initially null $\underline{T}^{+}_{\alpha}$ except for repercussions from the non-mentalistic outflow incited by $\underline{T}^{+}_{\alpha}$ in (65-1)'s short-term effects. (Thus, if (65-1) is $\underline{o}$'s perception of an armed robbery, and $\underline{T}^{+}_{\alpha}$ impels rash

actions leading to a bullet's penetrating the brain of a close successor of p, its long-term effects on cogitation in o's successors may well be severe.)]]

Moreover, our long skirmish with perceptual demonstratives recommends that canonical form 'o perceives that α is βish' for description of primary perceivings be elaborated by parsing the percept's subject-component α as an attributive-cum-target-marker whose most basic (primitive? prevalent? prototypic?) instances are concepts having the structure articulated in (65-1,2). To be sure, as progress in the study of mental mechanisms, this is important but scarcely astounding. For once we appreciate the prevailing syntactic complexity of nominals in singular propositions, the involvement of token-cues therein is simply one of many feature variations in propositionally structured ideation whose distinctive functional roles remain to be worked out. But this formulation does show how thought contents whose ordinary-language expressions cannot be freed of demonstratives are also describable without context-dependent use of demonstratives by a suitable theoretical vocabulary, thereby making it possible for us to subsume such mentation under communicable laws of thought. And emphasis upon forms (65-1,2) has major import for theories of mental representation. For it moves to center-stage the problem of reference by indefinite descriptions, and highlights the divide between representation by repeatable patterns and representation by nonrepeatable patterned events.

Even though the mental entities described by (65-1,2)'s content clauses are in both cases repeatable patterns of brain activity (or at least we have no good reason to suspect otherwise), a proposition that-a-T^{+}_{α}&χ-thing-is-βish whose token-cue T^{+}_{α} is non-null is semantically site-bound (contra site-free) in a way that that-a-χ-thing-is-βish is not.[55] For unless we abandon deictic semantics

[55] We shall here consider these to be distinct propositions even though I have suggested that what distinguishes them may be not conceptual content in a narrow sense but a fragment of modality contained in the first that is deleted from the second. Even if that proves to be the way to go, however, we shall still want semantic valuations not just of conceptual contents narrowly construed but of their modings as well.

altogether, we must hold that any proposition that-p^+ containing a non-null token-cue T^+_α does not in its own right stand for any factive object but only enables different occurrences ("tokens") of this pattern respectively to represent different events of a common kind at various locations picked out by each representation's own site under a denotation criterion selected by subpattern T^+_α.

Specifically, consider the representational contrast between (65-1) and (65-2), assuming that predicatives χ and β signify properties K and B, respectively, and letting here-and-now go heuristic proxy for any non-null token-cue T^+_α that might occur in (65-1). With their that-clauses understood to be genuinely of singular subject-predicate form, (65-1) and (65-2) each professes to describe some ϕ-moded representation of a kind-K object as having property B. But which kind-K object? If there exists exactly one K-thing, a_K, we can say for (65-2) that its a-χ-thing concept picks out a_K for repeatable proposition that-a-χ-thing-is-βish to represent as having B wherever this pattern recurs, veridically so iff object a_K in fact has property B. Whereas for (65-1), if there is exactly one K-thing a'_K in the immediate vicinity of cognizer-stage o, we can say that the event (or its t-core, or t-core locus) of o's-containing-activated-a-here-and-now-χ-thing-ideation refers to a'_K, with this particular occurrence in o of the that-a-here-and-now-χ-thing-is-βish proposition then representing a'_K as having B. Just where in space-time relative to o such an a'_K should be in order to qualify as the referent of this representation's nominal is no less vague than is commonsense intentionality's wont; but imprecision aside, the acceptability zone for a'_K is selected, under a function corresponding to the token-cue T^+_α for which we are taking here-and-now as proxy, by this particular representational event's location. That is, in sharp contrast to (65-2), what accomplishes representation in (62-1) cannot be just a repeatable thought pattern; it most include the where-and-when of some particular tokening of this pattern.

But what should we say is represented either by semantically site-free propositional pattern that-a-χ-thing-is-βish, or by the event of site-bound proposition that-a-here-and-now-χ-thing-is-βish being active in o, when the distribution of K-kind

is less ideal than just envisioned? If there are no Ks at all, then of course nothing is represented as having B either in (65-1) or in (65-2). But what if there exist many K-things or, for (65-1), more than one K-thing in the vicinity of o? In the latter case, token-cue T^{+}_{α} may well have the force of nearest-and-nowest; and even if not, repleteness of the χ-attributive can easily make negligible the probability of a local multiplicity of K-things given that one is present. (This is why everyday locutions of form 'this K-thing' seem unproblematically referential.) Yet no matter how richly detailed the χ-concept may be, if it is logically possible for the K-ness this purportedly signifies to recur, then we cannot plausibly presume that a-(somewhere-somewhen)-χ-thing picks out a unique referent given that it refers at all. So what does the repeatable proposition that-a-χ-thing-is-βish patternwise represent when there are many K-things? Two responses, both defiantly evasive, are appropriate here.

The first is to emphasize that multiple reference is indeed a major problem for semantical theory, one that arises far more pervasively than just from construing indefinite descriptions to function syntactically as nominals and which demands a fundamental reworking of standard philosophers' presuppositions about concept/object couplings. I have already spoken briefly to this issue on p. 257f., above, including the simplest (though not wholly adequate) way to handle it, and that is as far as we need to go here in sighting down new semantic-theoretical trails.

[Even so, this situation can be made somewhat more intelligible albeit no less disquieting as follows: The semantic status of proposition that-a-χ-thing-is-βish for a given cognizer-stage o may be viewed as essentially equivalent to that of theoretical proposition that-Johm-is-βish where 'Johm' is syntactically a proper name "implicitly" defined by o's acceptance of the minitheory consisting solely of the postulate that-Johm-is-a-χ. More loosely, the latter is simply the case where the Johm-nominal's role in o's conceptual economy makes the property signified by χ the sole criterion of Johm's identity—i.e.,

o's only grounds for judging whether this particular object is John is the strength of o's conviction that-this-is-a-χ. Virtually every proper noun in your real-life vocabulary--'Socrates', 'Boy George', 'Los Angeles', 'NATO', etc.--expresses a concept having this status for you; and even if you question that any class of disjoint entities picked out by what (say) 'Boy George' means to you should have cardinality greater than one, many different overlapping regions of space-time equally qualify for you as the location of this name's referent. (That remains true even if, contrary to reason, Kripke, Kaplan, and certain other modern philosophers of language are correct to hold that the referent of this name for you is determined not merely by your own understanding of it but by its history in your wider linguistic community.) You may shrug off such haze of referent locations as nothing more than the vagueness which, to one degree or another, invests all commonsense concepts-in-use; and you would be right not to be unduly exercised over it. (Imprecision works itself out as need arises, and gratuitous exactitute is generally counterproductive.) Yet philosophy of language has failed abjectly to give us an insightful account of what vague representations represent. Indeed, perhaps the main reason for philosophical semantics' shameful record in this matter is its reluctance to confront concepts whose objects cannot plausibly be identified simply by meta-linguistic applications of nominalization and disquotation.]]

Whatever accounts of multiple reference may prove to be technically tenable, these will surely tolerate our saying that the propositional pattern that-a-χ-thing-is-βish, if veridical, carries the information that some K-thing exists which also has B. Or to be really expedient, we can for now say that this veridical proposition represents the existential fact that $(\exists x)[K(x) \& B(x)]$, leaving for future ajudication whether it also represents certain particular K-things as having B.

Secondly, it is important to be clear that any answers a theory of aboutness may give to what objects are represented by what ideas under what circumstances

simply do not matter for a science of mental mechanisms, except insofar as pursuit of tidiness for such answers may bias our preferences for which molar properties of complex dynamic systems are to count as intentional. That representational aboutness is an epi-phenomenon having no import for the nature or causal functioning of thoughts in no way diminishes this matter's human importance: Use/mention interchanges (quotation and disquotation transformations) in our commonsense dealings with words, together with our deeply felt need to enhance our intellectual proficiencies by evaluative critiques (reasoned approvals/disapprovals) of our concept-economy's management, quite properly drives us to search for coherent theories of representation. But once the question-begging inadequacy of disquotational semantics for normative guidance becomes apparent, we can only hope that realistic accounts of aboutness can be extracted from a scientifically sophisticated grasp of how concepts mediate environmentally adaptive human achievements--after some such understanding becomes available. In short, put it this way: If we are entitled to feel confident of anything in semantics, it is that the distinctive representational character of any particular idea (thought, meaning, concept) μ for organism-stage o lies in the conceptual role played by μ for o. But what is that if not simply o's constellation of dispositions for μ and its compoundings with other ideas to participate in the causal processes that, at one level of molar abstraction or another, have made o's internal state and external situation what they are and what they will become? When we learn enough about the nomic regularities that cash out this "conceptual role" promissory note, the logic of aboutness will fall into our hands if not quite like a ripe apple then at least with only moderate tugging.

What might communicable depictions represent?

The argument just developed, that o's-seeing-that-α-is-βish is paradigmatically a more determinate token-cued o's-seeing-that-a-T^{+}_{α}&χ-thing-is-βish abstractively containing o's-seeing-that-a-χ-thing-is-βish, in principle rehabilitates the prospect of communication by inner-pictures. Our two primary variants

of conjectured depiction in o's-seeing-that-α-is-βish, you will recall, are (a) that this is simply o*'s-having-β wherein o's perceptive brain-region o* stands for a particular object a through certain structural properties F_α of o* and represents this a as having a property B signified by activity pattern β, or (b) what represents a as having B is o*'s joint property F_α-&-β, with F_α in itself, not its bearer o*, being what refers to a. Earlier we presumed that the inner-picture embodiment F_α of nominal concept α would have to include o*'s spatio-temporal location, or something tantamount to that, in order to pick out an a in o*'s vicinity. And were it not for one complication, that construal would still seem appropriate when α has composition α = a-T^+_α&χ-thing with T^+_α carrying the force of here-and-now. That is, we could say roughly that the depictive token-cue here simply is the location L(o*) of brain-region o*, while the a-χ-thing remainder of F_α is some complex P_χ of structural features (e.g. shape and size) that can recur in many different locations. Then inner-picture event o*'s-being-at-L(o*)-&-having-P_χ-&-β, which is a site-bound (tokenwise) representation of a as having B, also abstracts into o*'s-having-P_χ-&-β wherein P_χ-&-β is a repeatable condition that carries the information that a K-thing has B. And since this pattern can be reproduced throughout o's continuant social community--i.e., many brain regions in o's successors and their collegial contacts can share structural character P_χ along with the capability of activation β --depictive retention and communication of this existential information becomes routine in theory.

This simple account of communicable depiction cannot stand, however, without a major albeit easy shift in its treatment of token-cues. For, consider a process of long-term memory idealized by (64-1,2) with T^+_μ in (64-2) taken to be null. If the t-core of (64-1) is o*'s-being-at-L(o*)-&-having-P_χ-&-β with P_χ-&-β a repeatable pattern, we can easily imagine mechanisms for mnemonic reproduction of this property in a brain-region o^*_μ in successor o' of o. But o^*_μ too will have some location L(o^*_μ); so o^*_μ's containing the site-free depiction P_χ-&-β of a K-thing's having B

will be embedded in the event $\underline{o}^*_\mu$'s-being-at-$\underline{L}(\underline{o}^*_\mu)$-&-having-$\underline{P}_\chi$-&-$\beta$, which purports to be a site-bound representation (in all likelihood incorrect) of some $\underline{K}$-kind object in the vicinity of $\underline{o}'$ as having $\underline{B}$. To obviate this problem, however, we need only conjecture that inner-pictures embody here-and-now not in the referent-determining locations of depiction events so token-cued, but in some structural character $\underline{S}^+$ shared by some but not most depiction sites. (As will be apparent without my belaboring the point, $\underline{S}^+$ can be just one of many graded token-cue alternatives in depiction.) Suppose, for example, that brain stages have disjoint sectors corresponding roughly to certain open modes of thought--a perceptive sector, a short-term retentive sector, a long-term recollective sector, etc.--while $\underline{S}^+$ is the property of being perceptive, i.e. being wholly part of a perceptive brain sector. (Note that $\underline{o}^*$ has $\underline{S}^+$ iff each part of $\underline{o}^*$ has $\underline{S}^+$, as wanted for dense depiction--see p. 245, above. Note also that a depiction site's structural property of being short-term retentive might similarly embody nearly-here-and-now, and so on for other commonsensical deictic demonstratives.) Then if memory processes lead from $\underline{o}^*$'s-having-$\underline{P}_\chi$-&-β to $\underline{o}^*_\mu$'s-having-$\underline{P}_\chi$-&-β, where $\underline{o}^*$ is perceptive in observer $\underline{o}$ and $\underline{o}^*_\mu$ is long-term recollective in a successor $\underline{o}'$ of $\underline{o}$, the first of these site-free existence representations abstracts from $\underline{o}$'s site-bound representation of a here-and-now $\underline{K}$-thing as having $\underline{B}$, whereas the lack of $\underline{S}^+$-ness in long-term recollective regions allows $\underline{o}^*_\mu$'s state to embody recall in $\underline{o}'$ of the information that a $\underline{K}$-thing has $\underline{B}$ without concommitant imputation that any such thing is present to $\underline{o}'$.

The cogitive merits of inner-pictures vs. inner-sentences.

We have been so long at the issue of demonstratives (unless you took my advice to pass over that discussion) that I had best remind you that our broader aim here has been to illuminate the possible nature of percepts by comparing their inner-picture and inner-sentence models. It is time for a summary evaluation of that contrast.

It may well have occurred to you, as we drew out the force of conjecturing the t-core of o's-seeing-that-α-is-βish to be a depiction of some object a as having a property B, that the account of representation this gave us seemed considerably estranged from commonsensically conceived propositional thought. And that the latter is nondepictive is indeed a reasonable conclusion. But this conclusion should not be seized too quickly, especially not under supposition (evident in the writings of many cognition theorists) that if a mental representation is not sentence-like then it is perforce not propositional. Indeed, one large virtue of the inner-picture model is to make clear that there is nothing intrinsically objectionable in the hypothesis that images may constitute propositional thought; it is only certain unfoldings of this prospect that appear dubious. And we are thereby warned that the inner-sentence construal of propositional thought is equally conjectural.

The prospect that percepts might be inner-pictures is badly tarnished by one roundhouse objection which, however, incurs considerable backlash against the glibness over predicate signification for which fn. 49 (p. 241) requested your indulgence. Suppose that when I ask you what you are holding, you see that-this-thing-in-hand-is-a-book. According to the *Principle of Dense Depiction (p. 245), in order for your perceiving's t-core to be an inner-picture o*'s-having-β (or o*'s-having-F_α-&-β), your brain-region o*'s activity pattern β must signify a Bookness property which abstracts from properties (including relations) of this book's mereological parts in isomorphism to the constitution of β by properties of o*'s subregions. But Bookness does not seem to be at all that sort of property. For surely an essential facet of what it is to be a book is to play a certain role in social communication; and whatever the details of that role, commonsense protests against its being abstractable just from properties of the parts in any partitions of objects we take to be "books." To be sure, commonsense may be ingenuously wrong about this. For if the "social role" included in Bookness is no more than a set

of dispositions to interact with readers in certain ways, then a book's disposition to function bookishly should indeed be a-derivative from configurations of book-part properties. Yet even if that is so, it seems scarcely conceivable that the manner of that derivation could be isomorphic to the composition of any molar brain-region activity β. Nor need we reach for social roles to make this point: One major ingredient of Bookness is the property of comprising an unspecified but appreciable number of layers (pages) that are disposed to preserve their individual physical integrities when jostled, but to separate freely save at one edge. The property of comprising n pages (n a determinate integer) can straightforwardly even if demandingly be represented in the structure of β; but how these pages' movement-dispositions might then also be represented depictively in the subregional activity patterns from which β abstracts boggles the imagination.

Were it not for one large demurrer, this line of argument would pretty thoroughly sunder inner-pictures from commonsense percepts. For with only routine adaptations it can be repeated for nearly any predicative concept we are able to verbalize when reporting what we see in natural settings. But do ordinary-language predicates in fact correspond to objective properties in the simplistic fashion we have been presuming? Is there really any Bookness out there for percepts to represent? That our book-concept is flagrantly vague is not decisive; for while we surely do not want an ontology that admits a fuzzy Bookness de re corresponding exactly to our fuzzy notion of this, we can hold instead that the latter signifies loosely, to a certain degree, each complex external property that would be signified by some ideal precisification of this concept. Yet suppose that book were already ideally precise for us in such fashion that its representational tie to the world could be explicated by a statement of form 'Anything x is a book if and only if S(x)', where 'S(_)' is a complex predicate that articulates the composition of Bookness. (E.g., 'S(x)' might begin, 'x consists of at least three separable but individually cohesive layers flexibly joined at one edge ...'.) Considering the opulence of

logical quantifiers, modal operators, and physical/social descriptive terms that would appear in '$\underline{S}(\_)$' (don't ask--it's worse than you think), we can easily conclude that even if there really is an objective Bookness signified with precision by '$\underline{S}(\_)$', its occurrences cannot possibly be depicted--especially not in any seeing-that-α-is-a-book, insomuch as the book-concept therein contains little of the structure that would be manifest in '$\underline{S}(\_)$'. (Whereas you can see-that-this-is-a-book with scarcely any mental load, the demands of seeing-that-$\underline{S}$(this) should far exceed your human capabilities. Indeed, it is a major challenge for any account of predicate signification to explain how seeing-that-α-is-a-book can have the same factive object as the unmanageable seeing-that-$\underline{S}(\alpha)$.)

Unhappily, this argument carries farther than one might wish. For not merely does it discourage conjecturing commonsense percepts to be depictions satisfying PDD, its discomfort with simplistic predicate semantics (and hence, in light of our analysis of demonstratives, with nominal reference as well) suggests more sweepingly that seeking insight into the nature of mental contents foremostly in terms of what, representationally, these are objectively about is a mug's game. Taking that rejection seriously needs not extinguish psychonomic concern for representation; but it does urge us to shift talk about what some given feature Υ of thought in fact represents to talk about what Υ purports to represent. That is, instead of claiming/conjecturing that Υ does represent an entity such that ..., we do better to say only that the psychonomic functioning of Υ in the mental system at issue is as though Υ represents something such that ... . And this "as though" appraisal is then to be cashed out in some account of principles that govern the behavior of whatever molar properties of cognizant systems fall under the categories of commonsense intentionality-talk, at least insofar as the latter do indeed rough in phenomena worth scientific study.

For example, this shift of focus would rewrite PDD to say merely, in essence, that if the t-core o*'s-having-β of o's-seeing-that-α-is-βish is a depiction,

then any event $\underline{o}^*_j$'s-having-β_j from which the former is a-derivative is the t-core of a seeing-that-α_j-is-β_j by $\underline{o}$ for some nominal concept α_j. The weakened PDD no longer requires any such $\underline{o}^*_j$'s-having-β_j to be a veridical representation whenever $\underline{o}^*$'s-having-β is, nor even to be about anything at all. Instead, it acquires psychonomic force from non-representational *principles idealizable by schema

$\underline{L}_{cog}$: For any observer $\underline{o}$, nominal concept $\underline{x}$, and predicate concept $\underline{y}$, if $\underline{o}$ perceives that-$\underline{x}$-is-$\underline{y}$ish, then $\underline{Result}(\underline{o},\underline{x},\underline{y})$,

wherein consequent clause $\underline{Result}(\underline{o},\underline{x},\underline{y})$ is in all likelihood a complicated conditional that needs supplementation by additional states of $\underline{o}$--e.g., a match of idea $\underline{x}$ or $\underline{y}$ to $\underline{o}$'s other active thoughts or dispositions thereto--for any mental occurrence to follow. . . .

[[For example, '$\underline{Result}(\underline{o},\underline{x},\underline{y})$' might include 'For any concept $\underline{z}$, if $\underline{o}$ actively or dispositionally believes that-all-$\underline{y}$s-are-$\underline{z}$s, then $\underline{f}(\underline{o})$ actively believes that $\underline{x}$ is a $\underline{z}$', where $\underline{f}(\underline{o})$ is a short-lag successor of $\underline{o}$. Then under $\underline{L}_{cog}$, $\underline{o}$'s-seeing-that-α-is-βish brings about $\underline{f}(\underline{o})$'s actively believing that-α-is-a-γ conditional on $\underline{o}$'s having the belief (either stored or activated) that-all-βs-are-γs.]]

Given $\underline{o}^*$'s-having-β, with $\{\underline{o}^*_j$'s-having-$\beta_j\}$ any collection of its abstraction-base fragments as just described, what then follows from $\underline{L}_{cog}$ under the weakened PDD is not just $\underline{Result}(\underline{o},\alpha,\beta)$ but all of $\{\underline{Result}(\underline{o},\alpha_j,\beta_j)\}$ for $\underline{o}^*$'s parts $\{\underline{o}^*_j\}$. Although production of occurrences from this array of Result-consequences depends on release of conditionalities in the latter by $\underline{o}$'s standing on other state dimensions, you can easily see how the depictive construal of $\underline{o}$'s-seeing-that-α-is-βish is far more jeopardous of system seize-up under $\underline{L}_{cog}$--that is, implication of conjoint occurrences which are in fact competetive--than would be a model of perception that does not require this perceiving to be concommitant with enormously many other perceivings by $\underline{o}$. What is dubious in this is not that there might be

laws under which some complex brain states with picture-like organization give rise to others. Rather, it is our prospects for formulating principles that apply indiscriminatively (within limits) to all abstractions from all parts of such brain-state configurations, and do so moreover in rough agreement with commonsense expectations about the flow of ideation.

Even so, it is conceivable that a sophisticated theory of depiction might tenably constrain PDD tightly enough to turn aside this seize-up threat. If we ignore PDD and its attenuations altogether, do any significant differences remain between inner-pictures and inner-sentences? There do indeed, even though waiving appeal to objective reference and some PDD-type opulence of nested representations there is little to identify a thought as "depictive" except deficiency in the distinctive features of inner-sentences. The bottom line of this balance sheet is simply that ordinary-language descriptions of perceptual contents impute these to have certain system properties which are not adequately realized by inner-pictures. So either those alleged properties are psychonomically spurious or whatever embodies them in brain processes is not depictive.

Consider again our primary perceptual format, seeing-that-α-is-βish. In most real-life instances of this (arguably, all without exception) subject-concept α is a target-marking nominalization $\alpha = \underline{T}_\alpha\&\chi$ of some attributive concept χ, as variously illustrated by everyday locutions 'this χ-thing', 'the χ-thing', 'a here-and-now χ-thing', 'χ-ness', and still others noted earlier. (Recall that even when we verbalize α by a bare demonstrative, lacking an explicit χ, the demonstrative generally functions in part--though only in part--to signal an attributive that we cannot readily express.) So if o's brain-region o* is the site of o*'s-seeing-that-$\underline{T}_\alpha\&\chi$-is-$\beta$ish, o*'s totality of properties must include some embodiment of the χ-concept as well as one of β. And however those embodiments are accomplished, they should not merely be **noncompetitive**--i.e., o*'s having the one must not preclude its having the other--but to be systematically so in that most options

for $\mathcal{K}$ exclude few if any options for β. One way to achieve this--the inner-picture way--is for the total-state space over sites of mental representations to factor into logically independent subspaces $\underset{m}{P}$ and $\underset{m}{Q}$ of two nomically distinct kinds of properties such that the options for $\mathcal{K}$ are abstracted just from $\underset{m}{P}$ while those for β abstract just from $\underset{m}{Q}$. (This is essentially all that we have done with the structural/activational distinction in setting up our two idealized percept models, though we have further anticipated that structural states $\underset{m}{P}$ should play a role in mental processes rather different from that of activational states $\underset{m}{Q}$, and have conceded that the abstracting of molar activation patterns may also draw upon structure in such fashion that certain options for $\underline{o}^*$'s $\underset{m}{P}$-state do indeed preclude some alternatives for β.) In contrast, the inner-sentence way to insure co-realizability of $\mathcal{K}$ and β is for $\underline{o}^*$'s mereological parts to include two disjoint subregions $\underline{o}^*_a$ and $\underline{o}^*_b$ whose activation states abstract into patterns $\mathcal{K}$ and β, respectively. Then there can be no conflict between $\underline{o}^*$'s part $\underline{o}^*_a$ having $\mathcal{K}$ and $\underline{o}^*$'s part $\underline{o}^*_b$ having β, even though causal antecedents that co-produce these compatible events might be hard to come by.

However, getting $\underline{o}^*$'s total state to contain $\mathcal{K}$ and β jointly does not suffice to make $\underline{o}^*$ the site of a seeing that-$\underline{T}_\alpha$&$\mathcal{K}$-is-βish. Ordinary language implies that $\underline{o}^*$'s state must also contain a target-marker $\underline{T}_\alpha$ tied to $\underline{o}^*$'s embodiment of $\mathcal{K}$ in some functionally significant fashion that does not similarly tie it to β. (Such a $\underline{T}_\alpha$ is needed, inter alia, to expand $\mathcal{K}$ into a nominal; and any model of propositional thought must further allow an array of $\underline{T}_\alpha$ alternatives, insomuch as ordinary language envisions many functionally different nominalizations of the same attributive $\mathcal{K}$.) At first thought, inner-pictures seem more adept at this than are inner-sentences: When $\underline{o}^*$'s embodiment of that-$\underline{T}_\alpha$&$\mathcal{K}$-is-βish is depictive, $\mathcal{K}$ is marked as nominal simply by being an abstraction from $\underline{o}^*$'s state in structural-property space $\underset{m}{P}$, while abstraction β from $\underline{o}^*$'s $\underset{m}{Q}$-state is the concept predicated of this nominal simply by virtue of its being of the activational (contra structural) kind. And to complete the inner-picture's nominal by adjoining $\underline{T}_\alpha$ to $\mathcal{K}$, we can

envision that structure space $\underset{m}{P}$ in turn factors into subspaces $\underset{m1}{P}$ and $\underset{m2}{P}$ such that choices for χ abstract just from $\underset{m1}{P}$ while $\underline{T}_\alpha$ and its alternatives abstract just from $\underset{m2}{P}$. (What then ties $\underline{T}_\alpha$ to χ rather than to β is their common status as structural.) In contrast, the inner-sentence model partitions $\underline{o}^*$'s subregion $\underline{o}^*_a$ containing the percept's nominal into two sub-subregions $\underline{o}^*_{a1}$ and $\underline{o}^*_{a2}$ whose activation states respectively abstract into patterns χ and $\underline{T}_\alpha$. That $\underline{T}_\alpha$ joins χ rather than β to form the percept's nominal, while β rather than χ is the percept's predicate, is accomplished for this inner-sentence by some antisymmetric complex of structural relations among $\underline{o}^*_{a1}$, $\underline{o}^*_{a2}$, and $\underline{o}^*_b$. (As a fanciful illustration, the structural format for predicating a concept in $\underline{o}^*_b$ of a concept in $\underline{o}^*_a$ might be $\underline{o}^*_a$'s being spatially surrounded by $\underline{o}^*_b$, like a fried egg's enclosure of its yolk by its white; while $\underline{o}^*_a$'s containing $\underline{T}_\alpha$ and χ as the unified subject of a monadic predication, rather than as a 2-tuple of nominals for a relational predication, might consist in $\underline{o}^*_{a1}$ and $\underline{o}^*_{a2}$ being spatially contiguous, unseparated by $\underline{o}^*_b$, like a double-yolked fried egg in contrast to two single-yolked eggs fried with whites run together.)

However, this inner-picture model of the subject/predicate distinction implies that the descriptive contents embodied in depictions divide inflexibly between structural concepts that are inherently nominal and activational ones that cannot be nominalized. Thus if seeing-that-this-χ-is-βish is depictive, its converse seeing-that-this-βish-thing-is-a-χ is constitutionally impossible. And inner-picturing of conjunctive predications is also dicy. For whereas an inner-sentence embodiment of $\underline{o}$'s-seeing-that-α-is-$\beta_1$ish-and-$\beta_2$ish puts its two predicate concepts into disjoint brain sites $\underline{o}^*_{b1}$ and $\underline{o}^*_{b2}$ and so risks no incompatibility between the β_1-pattern in $\underline{o}^*_{b1}$ and the β_2-pattern in $\underline{o}^*_{b2}$ regardless of what those may be, an inner-picturing of this proposition must co-exemplify activity patterns β_1 and β_2 in the same brain site $\underline{o}^*$, which is possible only if they do not compete. Of course, there should be many joint options for β_1 and β_2 that are indeed non-competitively realizable in the total state of $\underline{o}^*$, especially ones that are composed

PDD-wise in isomorphism to the composition of co-exemplifiable molar attributes of external objects. But as inner-picture theories lay claim to a repertoire of predicates increasingly beyond the at-best meager range of ones that can satisfy PDD, competition among those is bound to grow increasingly prevalent. (As discussed in Chapter 5, noncompetitive pattern variables are hard to come by.) To be sure, competition between β_1 and β_2 is no impediment to embodiment of o's-seeing-that-α-is-$\beta_1$ish-and-that-α-is-$\beta_2$ish as a pair of depictive events, o_1^*'s-having-F_α-and-β_1 and o_2^*'s-having-F_α-and-β_2, wherein o_1^* and o_2^* are disjoint brain regions of o that each have the repeatable structural condition F_α constituting nominal concept α. (For the end of a long story on F_α-repeatability in depiction, see p. 273f., above.) But that is not the same, either commonsensically or in all likelihood psychonomically, as a percept in which concepts β_1 and β_2 are each coupled predicatively with the very same embodiment of the α-nominal.

In short, the special format of inner-sentences gives these a compositional versatility that far outstrips the differential realizability by depiction of commonsensically conceived propositions. Yet for this to urge conclusion that percepts are more sentence-like than picture-like, it needs also to be argued that these propositional differences, cleanly distinguishable by inner-sentences but not reliably so if at all by inner-pictures, genuinely matter for cogitive information processing and moreover occur not just in nonsensuous ideation but in perception as well. In particular, given that

seeing that-this-χ-is-βish

is a realizable percept, is each of

seeing that-this-βish-thing-is-a-χ ,

seeing that-this-thing-is-a-βish-χ ,

seeing that-this-thing-is-βish while also
seeing that-this-thing-is-a-χ

likewise a realizable percept whose nomic force differs appreciably from that of the others in this almost-but-not-quite paraphrastic group? Commonsensically the answer seems clearly affirmative, at least regarding nomic import: Intuition tells me that my seeing that-this-βish-thing-is-a-χ, or my seeing that-this-thing-is-a-βish-χ, must be transformed into an awareness that-this-χ-is-βish before it can induce adjustments in the strengths of my generalized beliefs (dispositional or activated) that-all/most/many/few/scarcely-any/no-χs-are-βish, or before it can interact with my standing opinion that-most-βish-things-are-γs to make me suspect that-this-χ-is-a-γ. (To get clear on the point here, write out these percept/generality/conclusion triples as putative inference schemata.) And it seems doubtful that I could infer that-this-thing-is-probably-a-γ from my generalized belief that-most-βish-χs-are-γs and my observing both that-this-thing-is-βish and that-this-thing-is-a-χ were not my observations more tightly unified as an awareness that-this-thing-is-a-βish-γ.

To be sure, only an innocent or a philosopher of mind would think that folk psychology proffers articulate views on what grammatical forms of propositional consilience govern our real-life inferences. But the salient point here is simply that we have every reason for confidence that variations such as just noted in the syntax of logically equivalent propositions do indeed make considerable difference for how modings of these propositions function in cogitive processes, even if such effects are still poorly understood by psychonomic science. And everyday perception-talk affords no suggestion that the attributive concepts in perceptual nominals cannot be predicated. Not only is it commonplace for you to see that the book before you is blue, it is equally mundane that given a different preparatory set you might just as readily see instead that this blue thing before you is a book, or that this thing before you is a blue book.

So can we then conclude that percepts are not depictive? Not exactly. What does seem implied is the inadequacy of inner-pictures for modelling thoughts correctly

describable by that-clause completions of Psi-verbs, perceptual or otherwise. But our earlier Posit (p. 226) notwithstanding, it is far from clear how much of the perceptual-experience story is best told in these traditional terms. Beyond the problematic status of feelings and emotions (see fn. 23, p. 137a above), phenomenological research provides ample reason to question whether all conscious awareness is structured as moded propositions. (See especially Hunt, in press.) And as I gaze across my cluttered study through its broad rain-speckled window past the foreground birch and across the valley toward the freeway traffic, only a little amateurish introspection on my visual experience suffices for me to misdoubt that this bright, richly splendored vista with its me-in-the-world aura is a largish collection of disjoint concepts on display in separate little bins whose layout determines what is nominal to what predicate. Arguably, what is most consciously vivid in the phenomenology of perceiving is a pre-conceptual stage of sensuous erlebnis that blends introspectively into its most direct propositional consequences ("interpretations") but is not appropriately described in the same terms or grammatical format as the latter.

That the classic distinction between pre-conceptual sensation and perceptual judgment may still be cogent for the study of mind (so long as the having of sensations is not mistaken for a perceiving of them) is only incidental to the larger issue here: Insofar as the idiom of intentionality carries presumption that internal conditions so characterized have sentence-like compositions and inference-like effects, it is psychonomic folly to apply that idiom indiscriminately--as does modern "information processing" jargon--to all stages of central arousal from stimulus reception to motor outflow. There is simply no good reason to think that organisms pervasively work that way, even if some parts of the system at molarity levels creamed off by folk psychology may do so.

It would be fair but fatuous to conclude that although inner-pictures ultimately prove deficient to model the constitution of propositional perceiving

as folk psychology conceives of this, it is not thereby precluded that the preceding sensory stage of input has the iconic character suggested by modern research in, inter alia, the Sperling tradition. Though perhaps worth voicing, this broad-mindedness is largely vacuous insomuch as any partitionable object, in particular any brain-region o*, is a proto-picture in that its totality of attributes makes it depictive in all respects save possibly, like a work of nonrepresentational art, its failure to be a picture of anything. For o* has parts, and its parts have parts, all integrated by structural relations and embodying molecular parts-properties which abstract into nestings of more holistic features of o* and its subdivisions. And that is all a picture is in its own right; any full-blooded depicting it may accomplish is imposed by some extrinsic correspondence--cognitive, statistical-covariational, fortuitous isomorphism, or whatever--of its parts with components of something else, and properties of the former with properties of the latter, in a fashion satisfying some counterpart of PDD for this mapping. (Like all representations, depictions are relative to some rule of correspondence which, however, needn't be cognitive aboutness albeit that is prototypic.) Only when what we are saying about o*-like things picks out particular configurations of abstractions from these objects' total states as our target of concern does the question arise whether these selected complexes are structured more like pictures than like sentences. So the imageal status of pre-conceptual sensation (or of any other central-processing stage) cannot become a well-mounted issue until we regiment the events in question under descriptions that discriminate the conjecturedly salient aspects in which sensory arrays resemble and contrast with one another, and adduce some manner of viewing these as representations of something else, preferably as "information" (in the statistical sense) about conditions in the eliciting environment. As brought out by our Molar Photography heuristic in Chapter 5 (p. 203ff.), the gulf between that aspiration and its hard-science fulfillment remains awesome.