Predication and Cartographic Representation

Michael Rescorla

Abstract: I argue that maps do not feature predication, as analyzed by Frege and Tarski. I take as my foil (Casati and Varzi 1999), which attributes predication to maps. I argue that the details of Casati and Varzi's own semantics militate against this attribution. Casati and Varzi emphasize what I call the *Absence Intuition*: if a marker representing some property (such as mountainous terrain) appears on a map, then absence of that marker from a map coordinate signifies absence of the corresponding property from the corresponding location. Predication elicits nothing like the Absence Intuition. "$F(a)$" does not, in general, signify that objects other than a lack property F. On the basis of this asymmetry, I argue that attaching a marker to map coordinates is a different mode of semantic composition than attaching a predicate to a singular term.

§1. Linguistic and cartographic representation

Anyone can see that linguistic representation differs markedly from pictorial, diagrammatic, or cartographic representation. But what are the differences, and how deep do they run? Despite intense interest stretching over many decades, philosophers have found these questions remarkably difficult to answer.

Focusing on cartography, I will try to isolate a precise sense in which sentences and maps have different representational formats. Sentences and maps are complex representations whose components contribute systematically to their representational content. A theory of linguistic (or cartographic) representation must exhibit the compositional mechanisms through which a sentence's (or map's) components contribute to its content. My thesis is that linguistic and cartographic compositional mechanisms are fundamentally distinct. Specifically, I will argue that maps do not feature *predication*, perhaps the most basic linguistic compositional device. In advocating this thesis, I follow Sloman (1978) and Millikan (2004, p. 93). I diverge from Casati and Varzi (1999, p. 192), Head (1984), and Pratt (1993), who impute predicational structure to maps, and also from Sellars (1979), who analogizes predication to cartographic representation.

§2. Frege and Tarski on predication

Frege's analysis of logical form revolutionized the philosophical study of predication. Frege treated semantic composition as *functional application*. Specifically, he treated predicates as referring to functions from objects to truth-values. Building on this idea, he analyzed the deductive relations between an unprecedented range of sentences, including sentences expressing multiple generality. He also indicated how to derive truth-conditions for those sentences.

Frege did not codify his discussion of truth-conditions into a formal semantics. Moreover, the thesis that predicates refer to functions proved tendentious. See (Burge 2005, pp. 18-22, pp. 99-103) for helpful discussion. Tarski domesticated many of Frege's insights within a formal theory of truth while avoiding tendentious appeal to Fregean functions. On Tarski's approach, the distinctive semantic contribution of predicates is to be *true of* objects (in more Tarskian terminology, to be *satisfied* by objects). If "$F(x)$" is a one-place predicate, then, assuming "a" denotes some object *den*("a"),

(1) "$F(a)$" is true iff "$F(x)$" is true of *den*("a").

As Quine puts it, "[p]redication joins a general term and a singular term to form a sentence that is true or false according as the general term is true or false of the object, if any, to which the singular term refers" (1960, p. 96). Generally, the semantic role of an *n*-place relation symbol "$R(x_1,\ldots, x_n)$" is to be true of sequences of objects, and

(2) "$R(a_1,\ldots, a_n)$" is true iff "$R(x_1,\ldots, x_n)$" is true of (den("a_1"), …, den("a_n")).

In what follows, I construe "predicate" broadly to include *n*-place relation symbols.[1]

As Davidson (2005) emphasizes, (1) and (2) do not explicitly associate predicates with functions or any other entities. Philosophers continue to debate whether Fregean functions add any explanatory value to our semantics beyond what clauses like (1) and (2) express. See (Dummett, 1981 p. 166-181) and (Davidson 2005, p. 120-140) for opposing viewpoints. I will not address this debate here, nor will I explore the differences between Fregean and Tarskian semantics. Instead, I want to emphasize the following four similarities. First, just like functional expressions, predicates contain "argument-places", i.e. open positions that can be filled by terms. Second, to convert a predicate into a truth-evaluable formula, one must fill all its argument-places either with denoting terms or bound variables. Third, even though Tarski does not literally associate predicates with functions from objects to truth-values, the satisfaction relation *induces* a function from objects to truth-values. Fourth, even though Tarski does not treat predication as a literal instance of functional application, (1) preserves the intuitive idea that "$F(x)$" carries or maps den("a") into the truth-value of "$F(a)$". Given these four points, there is a clear sense in which Tarski retains a broadly Fregean treatment of predication as functional application. Accordingly, I will speak of the "Frege-Tarski" conception of predication.

Despite occasional dissent, the Frege-Tarski conception has dominated philosophical study of predication for the past century. Truth-conditional semanticists offer various theories to handle phenomena such as reference failure, modality, propositional attitude reports, vagueness, and context-sensitivity. Although sometimes departing dramatically from Frege and Tarski, these theories usually retain, either explicitly or implicitly, the functional conception of semantic composition. That conception figures explicitly in (Heim and Kratzer 1998) and implicitly in (Larson and Segal 1995), two recent textbooks on semantics of natural language that agree on little else. A few philosophers, such as Sellars (1979) and Strawson (1964), try to analyze predication without deploying anything like the Frege-Tarski conception. Davidson (2005, pp. 98-119) forcefully argues that such attempts do not satisfactorily handle the semantics of complex predicates and objectual quantifiers.

The thesis of this paper is that maps lack predicational structure, as analyzed by Frege and Tarski. Like sentences, maps contain denoting terms, i.e. terms that denote or purport to denote particular objects. Unlike sentences, maps do not feature a compositional mechanism whereby such terms fill the argument-places of a predicate that carries their denotations into a truth-value. Instead, maps feature a quite different compositional mechanism for combining denoting terms into truth-evaluable representations.[2] Sloman (1978, pp. 145-176) endorses a similar thesis. He attaches it to a general distinction between *analogue* representations, which purport to replicate relevant features of what they represent, and *Fregean* representations, which exhibit the function-argument structure isolated by Frege. Sloman does not argue for the crucial thesis that maps lack predicational structure. My goal is to motivate and defend that thesis.[3]

§3. Cartographic representation

Map semantics has been less widely studied than the semantics of natural or artificial languages. We lack a canonical cartographic semantics comparable to Tarskian semantics for the predicate calculus. But a few researchers have offered suggestions, including (Casati and Varzi

1999), (Davis 1986), (Dilworth 2002), (Lemon and Pratt 1999), (Leong 1994), (O'Brien and Opie 2004), (Pratt 1993), and (Sloman 1978).[4]

Just like a sentence, a map represents the world as being a certain way. Just like a sentence, the map is correct iff the world is as the map represents it as being. I will describe maps as "true" or "false" according as they correctly or incorrectly represent the world. Some philosophers may question this usage, citing the many differences between maps and sentences. But even these philosophers should accept that *some* notion of representational correctness applies to maps. These philosophers can replace my talk about "truth" with talk about "representational correctness" or "veridicality." Similarly, I will say that a map's components "denote" objects and properties in the physical world. Cartographic semantics must unveil the compositional mechanisms of maps, that is, the mechanisms through which elements with certain denotations compose to yield a map with a given truth-condition.

A map represents how various objects and properties are distributed in physical space. It does so by purporting to *replicate* relevant geometric features of the spatial region it represents, as Russell (1923, p. 152) already noted and as virtually all subsequent commentators have emphasized. City maps purport to replicate *metric* structure (e.g. angles, relative distances), whereas subway maps purport to replicate only *topological* structure (e.g. connectedness, adjacency, containment). As a first attempt at characterizing how geometric structure contributes to a map's representational content, we might offer the following rough description. A map is composed of *coordinates* and *markers*. The coordinates are endowed with geometric structure, and the markers are located at certain coordinates. Coordinates denote locations in the physical world, while markers denote physical objects or properties. A map is true iff:

(3a) The distribution of markers across coordinates replicates the distribution of denoted objects and properties in physical space; and

(3b) The geometric relations between coordinates replicate salient geometric relations between denoted locations in physical space.

One of my main concerns will be how to state (3a) more formally. I will not discuss (3b) in much detail. Presumably, however, a more formal statement would employ the notion of a structure-preserving function (i.e. *homomorphism*) from the map to the world. For instance, a more precise formulation of (3b) for metric maps might run as follows: for any coordinates a and b, $dist(a, b) = k \times dist(den(a), den(b))$, where k is a constant scaling factor, $dist(x, y)$ is the distance between x and y, and *den* carries items on the map to their denotations.

Philosophers sometimes claim that maps possess intrinsic, as opposed to conventional, representational significance. To use Peirce's terminology: maps are *iconic*, not *symbolic*. (Goodman 1976) and (Sloman 1985) argue that a map's representational content is no more intrinsic than a sentence's. Just like a sentence, a map is true only relative to a denotation relation. In general, denotation is just as conventional for a map as for a language.

What associates a map with a given denotation relation? For instance, what makes something a map *of* London? That is an extremely difficult question. It is a special case of the problem of intentionality, analogous to the question: what makes a word denote a given referent? These are profound questions. Luckily, we do not need to answer them here. We are studying the compositional mechanisms through which a map's components contribute to its content, not the meta-representational question of what it is for those components to possess denotations.

In what follows, I confine attention to extremely simplistic maps. For instance, I make no attempt to address the variety of projection mappings modern cartography employs to produce 2-

D maps of the 3-D globe.[5] Still, I hope my main conclusions apply not just to the simplistic examples discussed below but to the full range of maps employed in contemporary practice.

§4. The internal structure of maps

Should predication, as described by Frege and Tarski, figure in a formal map semantics?

Just as one can insert a singular term into an argument-place of a predicate, one can insert a marker into a coordinate of a map. So we might try regarding coordinates as predicate argument-places. For instance, consider the following map:

(M_1)

	1	2	3	4	5
x		*A*			*C*
y				*C*	*C*
z		*B*	*B*	*B*	

Each of the fifteen grids is a coordinate, corresponding to a grid in physical space. Each marker represents either an object or a property, such as mountainous terrain. Placing a marker inside a map grid signifies that the corresponding object or property occupies the corresponding grid in physical space. We might analyze map M_1 as composed of a fifteen-place relation, "$M(x_1, \ldots, z_5)$", true of a sequence of objects or properties only if the objects or properties are found at the appropriate locations. Since maps often contain infinitely many coordinates (Kitcher and Varzi 2000), a general development of this approach would presumably require infinitary logic.

On closer inspection, however, the analogy between argument-places and coordinates falters. In Fregean terminology, deleting markers from a map does not yield the kind of "unsaturated" representation that results when we delete singular terms from a sentence. If we delete "Jack" from "Jack is hungry," what remains is no longer evaluable as true or false. It does not represent the world as being a certain way. If we delete marker *A* from map M_1, what remains is evaluable as true or false. It still represents the world as being a certain way. If we delete both *A* and x_2, then the result is a map with a hole in it, but nevertheless a map evaluable as true or false. Thus, a map *sans* marker differs dramatically from a sentence *sans* singular term, i.e. a predicate. We should not treat M_1 as composed of a fifteen-place predicate, because such a predicate attains a truth-value only if all fifteen argument-places are occupied by denoting terms or bound variables, whereas M_1 attains a truth-value even if certain coordinates are unoccupied.

Although there are various ways one might respond to this argument, the analogy between coordinates and argument-places seems sufficiently dubious that I will focus attention elsewhere. Specifically, I will focus on the suggestion that markers are predicates. In the special case where a marker denotes a particular object, we can treat it as denoting a property that only a single object instantiates, much as "socratizes" is true only of Socrates. We can then analogize placing a marker at a coordinate to predicating the relevant property of the relevant location. This proposal improves upon its predecessor, since an isolated marker, like a predicate, does not represent the world as being a certain way.[6]

To explore the proposal more systematically, let us examine the account from (Casati and Varzi 1999). Casati and Varzi confine attention to a very simple formal syntax that exclusively employs color patches as markers. Each color patch denotes some property. As Casati and Varzi put it: a "colored map region is like a sentence: it says that a certain individual (a spatial region) has a certain property. Like a sentence, it can be true (if the region is correctly colored) or false (if it is not)" (p. 191). Suppose color *C* appears on the map. Let $\sigma_{1,C}$ be the mereological fusion of all *C*-colored regions, and let $\sigma_{2,C}$ be its relative complement in the map. Consider the map in which $\sigma_{1,C}$ is colored C and $\sigma_{2,C}$ is uncolored; call it an "atomic map stage" of the original map. Let S_C denote the atomic map stage induced by C. We can now say:

(4) S_C is true iff *den*($\sigma_{1,C}$) has property *den*(C) and *den*($\sigma_{2,C}$) does not have property *den*(C).

(5) A map is true iff all its atomic map stages are true.

To Casati and Varzi, (4) and (5) suggest that "maps are propositionally structured --- albeit in a peculiar way" (p. 191). For (4) and (5) apparently provide recursive truth-conditions analogous to those of a Tarskian truth-theory.

In effect, (4) and (5) are more rigorous statements of (3a). What about (3b)? Casati and Varzi build geometric accuracy into the definition of "assignment function," i.e. the denotation function *den*. They countenance only those assignment functions that preserve relevant geometric (and mereological) structure. In this way, Casati and Varzi supplement (4) and (5) with a more formal statement of (3b). I will henceforth leave this supplementary demand unstated, but I intend it as implicit in all the semantic theories for maps that I discuss.

Does the semantics provided by Casati and Varzi isolate a predicational compositional mechanism? Casati and Varzi certainly think so, writing that a "color patch is an unsaturated predicate, which gets saturated when it is juxtaposed to a map region" (p. 192). But have they interpreted their own semantics correctly? In next section, I argue that they have not.

§5. The Absence Intuition

(4) does not explicitly mention the Tarskian satisfaction relation or Fregean functions from objects to truth-values. Hence, it does not explicitly embrace the Frege-Tarski conception of predication. However, this might seem like a relatively superficial discrepancy, since we can translate (4)'s talk about a spatial region "having" a property into talk about the region satisfying

the corresponding predicate. In this vein, say that marker *C* is "true of" some region iff the region has the property *den*(*C*) intuitively signified by *C*. We can then say

(6) S_C is true iff *C* is true of *den*($\sigma_{1,C}$) and *C* is not true of *den*($\sigma_{2,C}$),

which entails (4). Does (6) reconcile the Frege-Tarski conception of predication with Casati and Varzi's cartographic semantics?

It does not. Frege and Tarski treat predicates as functions carrying objects into truth-values. Frege takes the metaphor with functions literally, while Tarski domesticates it within the theory of satisfaction. Even metaphorically, (6) does not depict *C* as carrying *den*($\sigma_{1,C}$) into the truth-value of atomic map stage S_C. According to (6), the truth-value of S_C depends not only on whether *C* is true of *den*($\sigma_{1,C}$) but also on whether *C* is true of *den*($\sigma_{2,C}$). Attaching a marker to a region σ_1 does not yield a map that is true or false according as the marker is true or false of the spatial region denoted by σ_1. So *C*'s semantic contribution is not encapsulated by a function that carries denoted spatial regions into truth-values. Such a function may exist in some platonic heaven, but it does not capture how *C* composes with other symbols to represent the world. Thus, (6) abandons Frege-Tarski predication.

In this way, Casati and Varzi's own semantics militates against the philosophical gloss they place upon it. My quarrel here is not with the semantics itself, which is arguably the most rigorous and penetrating yet offered for maps, but rather with the claim that this semantics isolates predicational elements within cartographic representation.

We might replace (6) with a clause more closely analogous to (1):

(7) S_C is true iff *C* is true of *den*($\sigma_{1,C}$),

which entails

(8) S_C is true iff *den*($\sigma_{1,\,C}$) has property *den*(*C*).

But (8) is not the intuitively correct truth-condition. As a counter-example, suppose that *den*(C) is instantiated somewhere within the region denoted by $\sigma_{2,C}$.

***Map M_2*:**

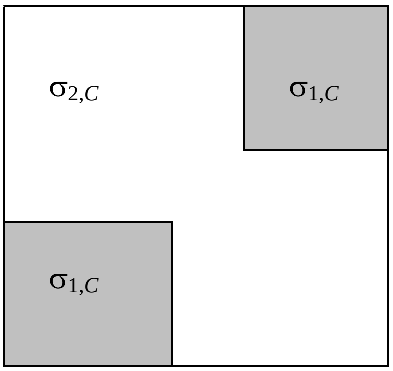

World

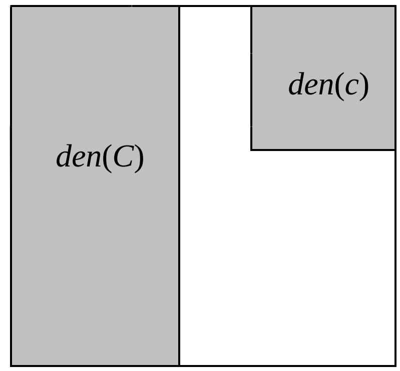

Intuitively, M_2 misrepresents the world. *Absence of C at a coordinate signifies absence of the corresponding property at the corresponding location.* Yet (8) judges M_2 true. In contrast, (4) judges M_2 false. Of course, *den*(C) might happen to be confined to *den*($\sigma_{1,C}$), in which case (8) would be true. Even so, there is a possible scenario in which M_2 is false but (8) predicts it is true. Thus, (8) fails the basic test of any semantics: it yields the wrong truth-condition.

Admittedly, there are some maps for which (8) yields the correct truth-condition. Consider a map depicting an ocean, using blue to depict water and green to depict scattered islands. Now suppose we cut holes to remove all the green. The resulting map contains blue at every remaining coordinate. (4) and (8) agree on this example. But that shows only that a false semantics can, in certain circumstances, yield correct truth-conditions. To decide between (4) and (8), we must examine cases where they disagree, not cases where they agree. When they disagree, (4) rather than (8) yields the correct truth-condition.

My argument, although centered around Casati and Varzi, does not exploit idiosyncratic details of their account. It exploits the following intuition: if a marker appears on a map, then absence of the marker from some coordinate signifies absence of the corresponding property from the corresponding location. This bedrock intuition is emphasized by (Casati and Varzi, 1999), (Hayes, 1985), (Lemon and Pratt, 1999), (Pratt, 1993), and (Schlichtmann, 1979). Let us call it the *Absence Intuition*.[7] Nothing like the Absence Intuition arises for ordinary predication. "$F(a)$" does not signify that objects other than a lack property F. That might sometimes be conversationally implicated by an assertion of "$F(a)$", but it is not typically part of the sentence's content. This contrast between cartographic and linguistic representation poses a challenge to anyone who holds that attaching a marker to map coordinates is the same mode of semantic composition as attaching a predicate to a singular term.

There are two main ways one might try to meet the challenge. One might contend that sufficiently artful employment of Frege-Tarski prediction yields the truth-conditions I ascribe to maps. Or one might question the truth-conditions I ascribe to maps.[8] I discuss the first line of attack in §6 and the second in §7.

§6. Deriving cartographic truth-conditions with standard predication

In this section, I critique various theories that concede the truth-conditions I ascribe to maps but try to derive those truth-conditions through Frege-Tarski predication. I do not claim that my arguments decisively refute the proposed analyses, let alone the many other predicational theories one might provide. But I will argue that such theories incur heavy costs without offering any evident benefits. They yield no explanatory advances, and they saddle us with awkward questions that otherwise do not arise.

§6.1 Restricting permissible marker semantic values

Certain natural language predicates include exclusivity in their meanings. For instance, if "*F*" means "is the tallest man in the world," then "*F*(*a*)" entails that "*F*" is true of no object save *den*("*a*"). Sometimes, then, predicating a property of an object signifies that other objects lack the property. So one might object that we can easily reconcile Frege-Tarski predication with the Absence Intuition. Certain predicates carry an implication of exclusivity, while others do not. Map markers merely fall in the former category.

In evaluating this objection, we must recall a crucial asymmetry between map markers and natural language predicates. Predicates like "is the tallest man in the world" are exceptional. In general, predication includes no claim to uniqueness or completeness. In contrast, if the truth-conditions I ascribe to maps are correct, then attaching a marker to certain map coordinates but not others includes a claim to exclusivity as part of its representational import. Thus, the Absence Intuition is intrinsic to cartographic representation but not to ordinary predication. This asymmetry does not conclusively show that cartographic compositional mechanisms are non-predicational. One might suggest that the Absence Intuition derives not from cartographic compositions mechanisms, but rather from restrictions on permissible marker semantic values. Specifically, one might suggest that map markers are restricted to semantic values carrying an implication of exclusivity. But I will now argue that this suggestion renders mysterious the very fact that it is intended to accommodate: namely, that *all* maps elicit the Absence Intuition.

How might one develop the suggestion in more detail? Two proposals seem natural. First, we might propose that the property *den*(*C*) denoted by marker *C* always carries an implication of exclusivity. For instance, instead of taking *den*(*C*) to be the property *being occupied by trees*, we might take it to be the property *being occupied by trees*, *and being such that all other locations* *depicted by the map are not occupied by trees*. Second, we might revise the conditions under which marker C is "true of" a region. In §5, I suggested:

(9) C is true of R iff R has property $den(C)$.

One might instead propose:

(10) C is true of R iff R has property $den(C)$ and $den(M/R)$ does not have $den(C)$,

where $den(M/R)$ is the set of locations denoted by map coordinates not in R. (7) and (10) yield the desired truth-condition (4) for an atomic map stage. I focus here upon (10), although my critique also applies to theories that retain (9) while restricting possible values of $den(C)$.

An account based on (5), (7), and (10) imputes Fregean function-argument structure to cartographic semantics. Yet it retains a principled distinction between markers and predicates. The distinction, which my account locates in compositional mechanisms, now resides in the requirement that the satisfaction relation conform to (10). That requirement is alien to Fregean and Tarskian semantics. It raises a pressing question: why can markers assume only semantic values conforming to (10)? Consider M_2. A proposal based on (10) grants that M_2 is false if $den(\sigma_{2,C})$ has $den(C)$. But the proposal does not explain why M_2 is confined to interpretations generating this truth-condition. Why must C fall under (10) instead of (9)?

An analogous question arises if we take (4) as primitive: why does (4), rather than (8), provide the correct truth-condition? The answer is that a map purports to replicate geometric features of the region it maps, including the distribution of salient properties. The map is true just in case it replicates those features. Replicating a property's spatial distribution requires replicating both where it is located and where it is *not* located. So a map is true only if it correctly divides space into the region occupied by a given property and the region unoccupied by that property. In other words, the map falls under (4). A representational system governed by (8) would be a *different* representational system than the one we seek to explicate.

Although the foregoing explanation may seem fairly trivial, it is far more satisfactory than any answer available once we embrace (7). The compositional mechanism described by (7) does not, by itself, entail (4). It entails (4) only when supplemented with something like (10). In general, predicates can express any functions from objects to truth-values, not merely functions conforming to (10). If markers are predicates, then why can they express only such a limited range of functions? If markers are associated with a satisfaction relation, why must the satisfaction relation conform to (10)? No answer seems forthcoming. Yet we require an answer, for otherwise (10) is an unexplained constraint upon a more general representational system. Lacking an adequate answer, an account based upon (5), (7), and (10) leaves us baffled as to why the Absence Intuition arises for all maps. We should not settle for a theory that predicts *of any given* marker that it elicits the Absence Intuition while rendering mysterious the fact that the Absence Intuition arises *for all markers*. Brute stipulation of (10), unaccompanied by supporting elucidations, creates an intolerable explanatory gap.

A predicate contributes to atomic sentences by attaching to singular terms whose referents it maps into semantic values. A marker contributes to an atomic map stage by dividing it in a way that purportedly replicates how some property's distribution divides physical space. As (5), (7) and (10) illustrate, we *can* assimilate the latter representational function to the former. But the assimilation, rather than illuminating maps, forces them into a Procrustean mold that ill-suits them. The assimilation undercuts our ability to explain a crucial asymmetry between cartographic and linguistic representation: the Absence Intuition is intrinsic to one but not the other. Moreover, we accrue no compensatory benefits, since (4) and (5) seem quite satisfactory taken as primitive stipulations. I do not say that this argument decisively refutes an analysis based on (10). But the argument suggests that nothing motivates the analysis save an antecedent commitment to discerning predicational mechanisms within cartographic representation.

Although I have focused on (10), analogous worries face any attempt at explaining the Absence Intuition through restrictions on permissible marker semantic values, rather than through compositional mechanisms. Thus, my argument shifts the burden of proof to those who favor the former explanatory strategy over the latter. We *can* treat markers as predicates that assume only certain semantic values. But why *should* we, given the high costs thereby incurred?

§6.2 Incompatible properties

Consider a map that uses green and blue to depict islands scattered in an ocean. Land and ocean are incompatible properties, so wherever the map indicates ocean we can infer absence of land. There is no need to say that absence of green indicates absence of land. We can instead say that presence of blue indicates presence of ocean, thereby indirectly indicating absence of land. More generally, one might argue, absence of marker C from coordinate x indicates absence of $den(C)$ only because the presence of some *other* marker D at x indicates presence of a property $den(D)$ incompatible with $den(C)$. We can therefore preserve the Absence Intuition while choosing (8) rather than (4) as the truth-condition for map stage S_C. There is an additional map stage S_D ensuring that the overall map has the correct truth-condition.

One problem with the proposed analysis is that a map need not explicitly place a marker at every coordinate. M_2 contains only a single marker C: the color grey. $\sigma_{2,C}$ is unoccupied blank space. M_2 does not explicitly represent $den(\sigma_{2,C})$ as instantiating any properties, except of course absence of $den(C)$. Thus, an analysis based upon incompatible properties does not predict that the Absence Intuition arises for $den(\sigma_{2,C})$.

In response, one might insist that placing no markers at a coordinate is tantamount to punching a hole through the coordinate. One might therefore suggest that we construe blank space not as *absence* of a marker but as *presence* of a "hidden" marker X with intended interpretation "absence of any property denoted by other markers appearing on the map." On this view, a map does not overtly display its syntactic structure any more than a sentence of natural language. In both cases, we must offer an abstract description of relevant structural features. As Frege's analysis of multiple generality demonstrates, those features may differ markedly from surface grammar. Researchers on natural language semantics respond by associating natural language sentences with a *formal language*, whose syntax may differ dramatically from surface syntax. Casati and Varzi propose that we proceed this way for maps as well: "define the notion of a *formal* map, which is in some sense recognizably similar to that of an ordinary map, and we provide a semantics for (some classes of) formal maps… It will then be possible to ask whether some features of this semantics can be used to describe the semantic structure of ordinary maps" (1999, p. 189). We can allow marked disparities between a concrete map's surface structure and the formal map associated with it. In particular, a formal map might contain a hidden marker X even though the associated concrete map contains no such marker.

Suppose, then, that a map has an additional atomic stage S_X corresponding to X, with the intended interpretation proposed above. Applying (7) to S_X yields:

(11) S_X is true iff, for all markers C on the map, $den(\sigma_{1,X})$ does not have $den(C)$.

Combining (5), (7), and (11) yields the intuitively correct truth-condition for M_2. More generally, (11) successfully handles the Absence Intuition as it arises for blank space on a map.

But a further problem remains. Given markers C and D on a map, there is no guarantee that $den(C)$ and $den(D)$ are incompatible. Yet the Absence Intuition persists for those regions occupied by *D* and unoccupied by *C*, and vice versa. This problem already arises within Casati and Varzi's framework, since there is no reason why two distinct color markers must denote incompatible properties. The problem emerges even more vividly when we recall that actual maps, unlike the idealized maps considered by Casati and Varzi, allow a single coordinate to contain more than one marker. Maps employ various devices to locate several markers at a single coordinate: textural variation imposed upon colors, superimposed symbols (e.g. a bridge symbol superimposed over a river symbol), and so on. Consider a simple example:

(M_3)

	1	2	3	4	5
x		*C D*	*C*		
y				*A D*	*D*
z			*A B*		

Intuitively, M_3 is true only if *den*(*C*) occupies no grid except *den*(x_2) and *den*(x_3). (11) does not yield this result. Since y_4, y_5, and z_3 are occupied by markers, they do not belong to $\sigma_{1,X}$. Thus, (11) yields an incorrect truth-condition that is neutral as to whether *den*(*C*) inhabits *den*(y_4), *den*(y_5), or *den*(z_3). To fix this problem, we might insist that *den*(*C*) be incompatible with *den*(*A*), *den*(*B*), and *den*(*D*). That maneuver would indeed entail that *den*(*C*) is absent from *den*(y_4), *den*(y_5), and *den*(z_3). But the maneuver faces two problems. First, there is no clear reason why a proper assignment of semantic values to M_3 must conform to it. Second, if M_3 is true, then *den*(*D*) is compatible with *den*(*C*), since M_3 depicts both properties as located at *den*(x_2).

Here is a slightly more realistic example to illustrate the same point:

(M_4)

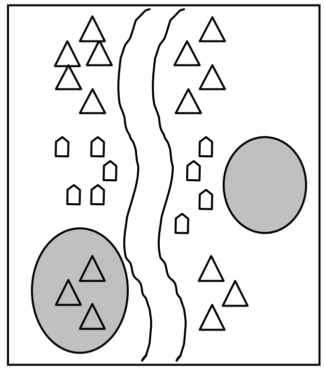

Triangles denote mountains; pentagons denote houses; the curved lines in the middle denote a river; grey denotes presence of some animal species. “Hidden” marker *X* cannot co-habit with grey, since *X* occurs only where other markers are absent. So (11) fails to predict that, if M_4 is true, then the right-hand ellipse occupied by the animal species is non-mountainous. Yet that is our understanding of M_4’s truth-condition. We cannot rectify this problem by urging that *den*(Δ) and *den*(*grey*) are incompatible. M_4, in its lower-left quadrant, explicitly represents the animal species occupying mountainous terrain.

In response to my argument, one might insist that, for each marker *C*, $\sigma_{2,C}$ is occupied by a hidden marker *X(C)* denoting absence of *den*(*C*). For instance, one might insist that M_4 contains a marker *X(*Δ*)*, present wherever Δ is absent, denoting absence of mountains. We would also need to posit an analogous marker *X*(⌂). It is not clear why we should clutter our theory with so many hidden markers just to preserve (8). Moreover, the proposal encounters a version of our worry from §6.1. When we apply a predicate to certain denoting terms, there is normally no requirement that we simultaneously apply an inconsistent predicate to all remaining denoting terms in our language. Why does an analogous requirement govern markers? If applying a marker to map coordinates is the same semantic operation as applying a predicate to a singular term, why does it fall under such special restrictions in the cartographic case? The proposal depicts basic cartographic phenomena as *ad hoc* impositions on a more general compositional mechanism, and it offers no compensatory benefits for the attendant loss in explanatory power.

In general, then, absence of marker C does not guarantee presence of a further marker D denoting an incompatible property.

§6.3 Markers as relation symbols

Another idea is to construe C as a *two-place* relation true of (σ_1, σ_2) just in case *den*(C) is located at σ_1 but not at σ_2. More precisely, we can say that a formal map consists of atomic stages that are atomic sentences: "$\Phi_C(\sigma_{1,C}, \sigma_{2,C})$", where $\sigma_{1,C}$ is the region occupied by C and $\sigma_{2,C}$ its complement in the map. We can then say:

(12) "$\Phi_C(x, y)$" is true of (R, S) iff *den*(R) has *den*(C) and *den*(S) does not have *den*(C).

When coupled with (5) and a suitable analogue of (2), (12) yields the correct truth-condition for the map as a whole. (12) implicitly invokes Fregean function-argument structure, since it treats a marker as inducing a function from ordered pairs of spatial regions to truth-values.

By analyzing maps through the framework of function-argument structure, (12) prompts us to ask why that framework's usual properties do not apply to maps. A version of our worry from §6.1 arises yet again: if markers are relation symbols, why must their semantic values conform to (12)? Another problem is that, if a proposal based on (12) were correct, one would expect that we could delete $\sigma_{2,C}$ from "$\Phi_C(\sigma_{1,C}, \sigma_{2,C})$", yielding "$\Phi_C(\sigma_{1,C},\ \ )$," an unsaturated representation that lacks a truth-value and that is true of S just in case *den*($\sigma_{1,C}$) has *den*(C) but S does not. However, no familiar operation upon maps corresponds to this deletion. Consider:

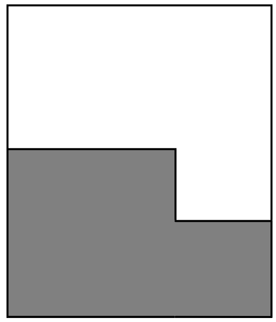

Let grey be marker *C*. Removing $\sigma_{2,C}$ from the above map generates a new, smaller map entirely covered with *C*:

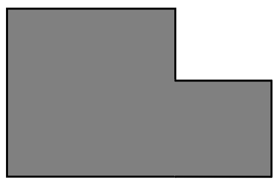

Unlike "$\Phi_C(\sigma_{1,C},\quad)$," this new map requires no supplementation to attain a truth-value.[9] Moreover, *C* by itself, unattached to any coordinates, cannot correspond to "$\Phi_C(\sigma_{1,C},\quad)$", since $\sigma_{1,C}$ makes no appearance in *C*. For similar reasons, no familiar operation on maps yields a representation akin to "$\Phi_C(\quad, \sigma_{2,C})$". If markers are relation symbols, why can't we form unsaturated representations like those we can form with relation symbols? No obvious answer is forthcoming, save brute stipulation that such representations are not syntactically well-formed. In contrast, a semantics that takes (4) as primitive does not encounter these questions, since it does not traffic in relation symbols.[10]

§6.4 Pratt on circumspection

The Absence Intuition may remind some readers of *circumspection*, a popular tool within non-monotonic logic (Brachman and Levesque 2004). The idea is to privilege models that assign certain predicates an extension "as small as possible." We define a reflexive, transitive ordering

$\leq_{P,Z}$ over the models of a theory, where P is a set of unary predicates and Z is a set of arbitrary predicates. If M and N are models, then $M \leq_{P,Z} N$ iff M and N assign the same extension to each member of Z, and the extension under M of each member of P is a subset of its extension under N. A set of formulas S *default entails* a formula T relative to $\leq_{P,Z}$ iff T is true in all models of S that are minimal relative to $\leq_{P,Z}$. This definition reflects the idea that one is default entitled to presume that certain predicates' extensions are as small as possible given what one knows.

Although circumspection might seem to resemble Casati and Varzi's semantics for maps, there are crucial differences. Circumspection formalizes a default entitlement to *presume* that extensions are as small as possible. In contrast, (4) builds that presumption into a map's truth-conditions. Circumspection presupposes a Tarskian approach to predication, and it offers classical truth-conditions, such as: "$F(a)$" is true iff "$F(x)$" is true of *den*("*a*"). It employs a standard logical consequence relation validating the normal inference rules of classical logic. It merely *supplements* these familiar semantic notions with a "default entailment" relation governing defeasible inference. Thus, unlike (4), circumspection is irrelevant to truth-conditions.

Pratt (1993, p. 82) emphasizes the Absence Intuition. He acknowledges that it poses an obstacle to analyzing maps through Tarskian predication. Nevertheless, he suggests that we can impute predicational structure to maps by invoking circumspection. Pratt does not provide many details, rendering his proposal difficult to evaluate. Initially, however, it seems questionable. As just emphasized, circumspection employs normal Tarskian predication to yield the usual truth-conditions. So it is difficult to see how circumspection, in itself, helps achieve the result that both Pratt and I desire: truth-conditions in accord with (4) rather than (8). More generally, it seems unlikely that we can generate the desired truth-conditions through straightforward application of "default reasoning," as studied within non-monotonic logic. Formal theories of default reasoning retain the usual truth-conditions of predicate calculus sentences.

§7. Correctness and completeness

Is it really so clear that absence of marker C from a map coordinate signifies absence of *den*(C)? In other words, is it really so clear that we experience the Absence Intuition? One might insist that we distinguish between *incorrect* maps and *correct but incomplete* maps. For instance, isn't there a sense in which map M_2, while incomplete, is correct?

In providing a semantics for maps, we are elucidating a *pre-existing* representational system. I can stipulate, "Whenever I draw a picture that locates a marker at σ_1 but not σ_2, you should understand me as describing what is located at σ_1 while remaining neutral about σ_2." I can also stipulate: "Whenever I draw a picture that locates a marker at σ_1 but not σ_2, you should interpret me as saying that the price of tea in China has risen about a certain level." Obviously, the latter stipulation does not relate in any interesting way to the compositional mechanisms of ordinary cartographic representation. Our question is whether the former stipulation is likewise irrelevant to those mechanisms, or whether it registers one possible manifestation of them.

Here we must recall the paradigmatic reason we employ maps: navigation. As Sellars (1979, p. 109) puts it, "the essential feature of the functioning of a map *as*, in a primary sense, a map is its location in the conceptual space of practical reasoning concerning getting around in an environment." We want to determine where a destination is located *and* to determine which trajectories will avoid obstacles such as mountains, trees, and so on. In this context, absence of physical properties is just as important as presence. For instance, anyone who views M_4 as veridical will conclude that the region depicted by it contains a non-mountainous middle section flanked by two regions of mountainous terrain. Moreover, any sensible navigator would exploit this information when traveling from east to west, plotting a trajectory to pass through the non-mountainous terrain.

This example illustrates how the Absence Intuition informs normal navigational practice. We might adopt a different representational system, one to which the Absence Intuition does not apply. Call such representations *quasi-maps*. Whatever the virtues of quasi-maps, we would not exploit them to avoid obstacles in the way we exploit actual maps to avoid obstacles. The truth-conditions of maps are intimately connected to their use during navigational practice. We cannot alter those truth-conditions, replacing (4) with (8), without distorting the practice. So the choice between maps and quasi-maps is not just the choice between two equally meritorious representational systems. Maps, unlike quasi-maps, harmonize with actual navigational practice.

In response to my argument, critics might observe that we employ maps for many purposes besides navigation. For instance, journalists often use maps to represent the spatial distribution of economic and sociological properties, such as majority support for a political party. These maps are intended for informational rather than navigational purposes.

I respond that, although maps serve diverse purposes, their potential use within navigation is central to their status as maps. Part of what makes a representation a map, as opposed to some other mode of representing the same information, is that one can exploit it to locate desired objects or properties *and* to avoid undesirable objects or properties. For instance, if marker *C* represents majority support for some political party, then a traveler who wished to avoid members of that party and who believed the map to be veridical would plot a trajectory through spatial regions denoted by coordinates from which *C* is absent. Even though a given map may not be intended for navigational purposes, any genuine map, as part of its essential representational nature, can occupy a certain characteristic role within navigation. I have argued that this characteristic role reflects our implicit grasp of the Absence Intuition.

To bolster my conclusion, I now critique three ways one might argue for a conception of correct but incomplete maps. First, one might concede the Absence Intuition but explain it by positing two distinct types of representational correctness. Second, one might again concede the Absence Intuition, explaining it not semantically but as a pragmatic phenomenon akin to Gricean implicature. Third, one might contest the Absence Intuition.

§7.1 Two notions of representational correctness?

Leong (1994) offers a formal semantics for maps that distinguishes two notions of correctness: *accuracy* (correct but incomplete) and *precision* (correct and complete). "Precision" corresponds to the notion of correctness employed by Casati and Varzi. "Accuracy" is much less demanding. Rather than requiring that the map depict a property at a location *if and only if* the property occupies the location, accuracy demands that the map depict a property at a location *only if* the property occupies the location. M_2 is accurate but not precise. (8) and (5) provide accuracy-conditions, while (4) and (5) provide precision-conditions. Thus, accuracy-conditions might seem to secure a role for predication within cartographic semantics.

This approach is not satisfactory. Cartographic semantics must explain how a map's components contribute to its content, i.e. to how it represents the world as being. The map is correct iff the world is how the map represents it as being. We now ask: does a map that omits a marker from some coordinate thereby represent absence of the corresponding property? Either it does, or it does not. If it does, and if the property is present, then the map is incorrect. If it does not, then the map may be correct even if the property is present. There is no room here for equivocation. We cannot say, "Well, in one sense the map represents the world as being such that the property is absent, and in another it does not." This is doubletalk. A map cannot both represent the world as being a certain way and simultaneously fail to represent the world as being that way. Nor can a map represent the world as being some way while simultaneously representing the world as being a different way, unless the map contradicts itself, which most maps do not. Thus, a proper conception of cartographic semantics leaves no room for two types of representational correctness. Either a map correctly represents how the world is, or it does not.

The issue here is not whether *accuracy* and *precision* are clear, well-defined notions. They both are. The issue is *which* notion is relevant to evaluating whether the world is as the map represents it as being. If the map represents the world as being such that its precision-conditions are fulfilled, then accuracy-conditions are irrelevant to cartographic semantics (except insofar as they contribute to a compositional account of precision-conditions). If the map represents the world as being only such that its accuracy-conditions are fulfilled, then precision-conditions are irrelevant to cartographic semantics (except insofar as they contribute to a compositional account of accuracy-conditions).

Intuitions may persist that maps admit multiple standards of representational correctness. My approach can accommodate these intuitions as follows. An incomplete map is incorrect, but various of its sub-maps may be correct. If we punch holes in a map wherever it fails to register some property, then the resulting sub-map may be correct. So there are different ways of evaluating a map's correctness. Those different evaluations feature a uniform notion of representational correctness, as applied to different entities.

But can't we just stipulate that certain maps are to be evaluated in terms of accuracy-conditions, rather than precision-conditions? Of course. We can likewise stipulate that "John is tall" is true iff John is the only tall person in the world. This latter stipulation constitutes a radical revision of the ordinary predicational compositional mechanism. Similarly, I claim that the former stipulation constitutes a radical revision of ordinary cartographic compositional mechanisms. I have supported my analysis by adducing intuitions about truth-conditions and by highlighting constitutive connections between cartographic representation and navigational practice. So far, we have encountered no rival analysis that seems nearly as compelling.

§7.2 A pragmatic explanation

The basic idea behind the pragmatic strategy is that an incomplete map is misleading because of what one can reasonably infer from the map rather than what it explicitly represents. For instance, the pragmatic strategy grants that, when faced with M_4, we conclude that there is a non-mountainous region between two mountainous regions. A semantic account, such as (4), (10), or (12), holds that this is part of M_4's content. A pragmatic account holds that, while warranted, it is not part of M_4's content: M_4 is correct even if there are mountains where M_4 contains no instances of Δ. One might model this idea more formally through circumspection or some other tool from non-monotonic logic. However, I will leave the idea at an informal level.

If M_4 does not explicitly represent absence of mountains in its middle section, then why may one reasonably infer that mountains are absent? Borrowing from Grice (1989), one might reason as follows: in communicative situations, agents seek to convey as much relevant information as possible; since the map-maker omitted Δ from certain coordinates, he must not have felt confident that the corresponding region contains mountains; the likeliest explanation for why he lacked this confidence is that the region does not, in fact, contain mountains.

To undercut this pragmatic analysis, consider a new map M_4* formed by punching holes in M_4 in scattered coordinates in the middle of the map. We may stipulate that it is common knowledge that the map-maker himself punched these holes in the map. Intuitively, M_4* remains neutral about the deleted locations whereas M_4 represents them as non-mountainous. No doubt there are circumstances in which we would conclude that the locations deleted from M_4* are non-mountainous. However, assuming that the purveyor of M_4* does not explicitly advertise it as comprehensive regarding mountains, we regard this conclusion as far more tentative, provisional, and conjectural for M_4* than for M_4. If we discovered that there *were* mountains in the relevant locations, we would berate the creator of M_4 quite vociferously for producing a misleading map, but we would have to admit that, technically, the creator of M_4* deserved no comparable criticism. Can the pragmatic account explain this intuitive asymmetry? That seems doubtful. The pragmatic account should apply just as readily to M_4* as to M_4, since the two maps allegedly have the same truth-condition and since, in both cases, the map-maker deliberately declined to label certain regions as mountainous. Apparently, then, pragmatic considerations do not explain why the Absence Intuition arises so much more vigorously for M_4 than M_4*. In contrast, a semantic account readily explains the asymmetry, since it assigns the two maps different truth-conditions.

§7.3 Contesting the Absence Intuition

A third response to my argument questions whether absence of a marker must signify absence of its referent. In this vein, Davis writes that "[t]he fact that a map does not show an object does not mean that it is not there; maps are not obliged to show everything, or anything, in a particular area" (1986, p. 31). Viewed as a description of cartographic representation *in general*, Davis's assessment simply ignores strong intuitions evoked by maps like M_4, not to mention how we exploit such maps during navigation. But there are certain cases in which Davis's assessment might seem more apt than mine. I examine a few such cases.

§7.3.1 Changes in resolution

Consider a map generated by a computer program such as Google Maps. As we zoom in, more detail appears. Often, the extra detail concerns properties signified by the less detailed map. For instance, a more detailed city map might depict smaller roads while the less detailed one depicts only major roads. Does my analysis entail that the less detailed map is incorrect?

No. The crucial point here concerns which properties the markers on the less detailed map represent. In a city map that shows only major roads, we construe road-symbols as denoting roads meeting some criterion, such as width, number of lanes, and so on. Under this construal, the map purports to depict only roads satisfying the relevant criterion. The map is false if it omits roads satisfying this criterion, but it may be true even if it omits roads not satisfying the criterion. When we zoom in, the more detailed map now purports to depict roads satisfying some broader criterion, not merely roads satisfying the original narrow criterion.

§7.3.2 Extraneous blank space

Consider a variant upon M_4:

(M_5)

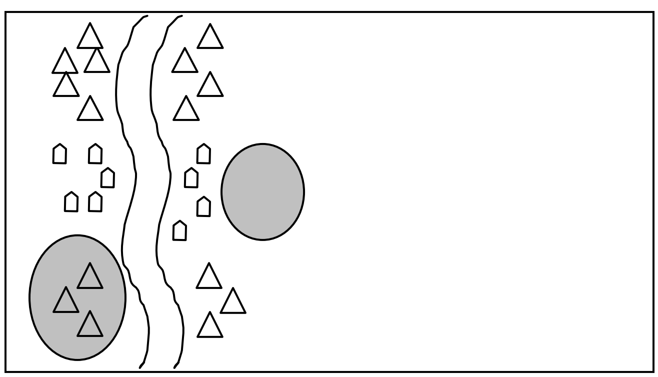

One might deny any robust intuition that M_5 represents absence of mountains or houses in the territory depicted by its right-hand side.

Several points deserve emphasis. First, even if we grant that the Absence Intuition does not arise for M_5's right-hand side, it arises quite vigorously for the left-hand side. Thus, the

Absence Intuition persists at a local, if not at a global, level. Second, if we place a single pentagon on M_5's right-hand side, then we experience a strong intuition that the revised map depicts absence of mountains on the right. Finally, our intuitions even for the original M_5 are unstable. One wants to say that it all depends on the map-maker's communicative intentions, the role played by M_5 in navigation, and so on. For instance, if it is common knowledge that the map-maker intends navigators to exploit M_5's right-hand side by relying on the belief that the corresponding territory lacks houses or mountains, then the Absence Intuition becomes quite compelling for the entire map. On the other hand, if M_5 is sketched hastily on a piece of paper much wider than necessary, and if its sole purpose is to depict terrain along the river, then the Absence Intuition becomes less compelling as we move rightward on the map.

We can accommodate these varying intuitions by invoking Casati and Varzi's distinction between formal and concrete maps. We can say that M_5 corresponds to infinitely many formal maps, which vary solely in terms of their right-hand boundary. That boundary might be located anywhere between the right-hand grey ellipse and the right-hand boundary of M_5 itself. Each formal map falls under the semantics given by (4) and (5). *Which* formal map is expressed by M_5 depends in complex ways upon M_5's history and its use during navigation. In many cases, those factors will not fix a unique formal map. Thus, a certain amount of vagueness informs the foundations of cartographic semantics. This vagueness is not a defect in my account. On the contrary, it is desirable, since it reflects our sense that, in many cases, there is no sharp boundary between where the Absence Intuition arises and where it does not. Since vagueness is pervasive in natural language, we should not be surprised that it arises also for cartographic semantics.

In this way, we can accommodate examples like M_5 without abandoning the Absence Intuition and without imputing predicative structure to maps.

§7.3.3 Terra incognita

Maps produced during the 17th and 18th centuries often employed the symbol "*terra incognita*" to signify ignorance. As a simplified version, consider:

(M_6)

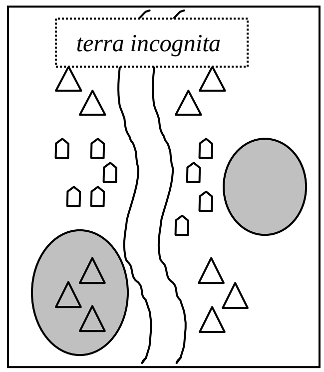

M_6 is true even if the area marked "*terra incognita*" contains mountains.

A crucial point here is that, in the region *not* denoted *terra incognita*, absence of Δ still signifies absence of mountains from the corresponding location. For instance, just like M_4, M_6 represents a non-mountainous region in its middle. Thus, the Absence Intuition persists, confined to those regions not denoted "*terra incognita*." A good theory should preserve elements of the simple account from §§4-5 while acknowledging the distinctive role played by "*terra incognita*." Such a theory will not attribute predication to maps. For the representational role it assigns to markers in regions not denoted "*terra incognita*" is fundamentally different from the role played by predicates in Tarskian semantics.

I propose the following analysis. The symbol "*terra incognita*" serves two representational functions. First, it denotes a property along the following lines: "terrain whose salient physical properties are unknown." Like any other marker, "*terra incognita*" signifies presence or absence of this property through presence or absence on the map. Second, "*terra incognita*" restricts the scope of each additional marker *C*. More precisely, let $\sigma_{1,TI}$ be the mereological fusion of the regions covered by "*terra incognita*", and let $\sigma_{2,TI}$ be its relative complement in the map. S_{TI} is the atomic map stage in which $\sigma_{1,TI}$ is covered by "*terra incognita*" and $\sigma_{2,TI}$ is blank. For all other markers C, let $\sigma_{2,C}/\sigma_{2,TI}$ be $\sigma_{2,C}$ minus $\sigma_{2,TI}$. S_C is now composed of $\sigma_{1,C}$, which is covered with C, and $\sigma_{2,C}/\sigma_{2,TI}$, which is blank. We now say:

(13) S_{TI} is true iff *den*($\sigma_{1,TI}$) is occupied by terrain whose salient physical properties are unknown and *den*($\sigma_{2,TI}$) by terrain whose salient physical properties are known.

(14) S_C is true iff *den*($\sigma_{1,C}$) has *den*(C) and *den*($\sigma_{2,C}/\sigma_{2,TI}$) does not have *den*($C$).

(5), (13), and (14) generate the intuitively correct truth-condition for maps containing "*terra incognita*," such as M_6. Thus, we can accommodate "*terra incognita*" with only minimal adjustments to the semantics from §4.

Andrews writes that "[b]lank space… may reasonably be interpreted as a negative judgment about the world, but it may just as reasonably be seen as a gesture of agnosticism" (1990, p. 15). Jean Baptiste D'Anville's 17th-century map of Europe and Africa left the African part mostly blank to indicate unknown terrain. One way to analyze such maps is to posit an implicit "*terra incognita*" located wherever blank space signifies ignorance. As in §7.3.2, a certain amount of vagueness will typically surround the boundary between regions where blank space signifies ignorance and regions where it does not. As a limiting case, we can imagine a map that contains an implicit "*terra incognita*" wherever blank space prevails. However, it is doubtful that many historical maps satisfy this description. Such a map is not very useful for navigation, since navigators usually care as much about which salient physical properties are absent as which are present. For instance, blank space in the European section of D'Anville's map is most plausibly interpreted as signifying absence of relevant properties.

The obvious objection to (14) is that it posits an *ad hoc* revision of (4). But it is not surprising that (4), which was prompted by reflection upon extremely simplistic examples, should require minor revision. Note that the predicative (10) and (12) fare no better in this regard, since they require, respectively, substituting $den((M/R)/\sigma_{1,TI})$ for $den(M/R)$ and $\sigma_{2,C}/\sigma_{2,TI}$ for $\sigma_{2,C}$. The pragmatic approach from §7.2 fares somewhat better. Rather than modifying the semantics, it can simply say that "*terra incognita*" cancels any implicature that $den(\sigma_{1,TI})$ lacks $den(C)$. While this contrast counts in favor the pragmatic approach, it does not seem sufficiently compelling to outweigh the serious difficulties raised in §7.2.

In short: although "*terra incognita*" requires us to temper our initial claim that absence of a marker *always* signifies absence of the corresponding property, it should not incline us to impute predicative structure to maps.[11]

§8. Implications of my discussion

My argument hinges upon a single intuitive distinction: attaching a marker to a set of map coordinates signifies that the corresponding property does *not* occupy certain locations; attaching a predicate to a singular term does not generally signify that any object lacks the property denoted by the predicate. At first blush, this distinction might not seem significant enough to bear much theoretical weight. I have argued otherwise. As we saw in §6, attempts to preserve the distinction while imputing predicative structure to maps raise awkward questions without offering compensatory benefits. As we saw in §7, attempts to eliminate the distinction by revising cartographic truth-conditions flout both intuition and navigational practice. Apparently, then, the distinction reflects fundamental representational disparities between attaching a marker to a coordinate and attaching a predicate to a singular term.

So what? Who cares whether maps feature predication? To address this question, I mention four broader implications of my discussion.

One implication concerns the proper direction of future research on cartographic semantics. Elucidating the semantic role of markers is essential to understanding how maps represent. Just as attaching a predicate to a singular term is the most elementary act of linguistic representation, attaching a marker to a coordinate is the most elementary act of cartographic representation. I have argued that cartographic semantics should not assimilate the latter operation to the former. As long as we do so, we are apt to pursue misguided theories like those critiqued in §§6-7. We should base cartographic semantics upon a fundamentally different foundation than sentential semantics.

A second implication concerns the question with which we began: how do cartographic and linguistic representation differ, and how deep do the differences run? A central problem in understanding representation is to explain how certain forms of representation (e.g. pictures, diagrams, or maps) differ in a principled way from sentences. My discussion isolates a precise sense in which maps and sentences differ profoundly in their semantic articulation: maps lack the function-argument predicative structure central to virtually all modern treatments of sentential semantics. In this way, my account applies formal semantical tools to the traditional philosophical enterprise of classifying representational systems.

Thirdly, my discussion bears upon the murky problem, introduced by Plato, sometimes called "the unity of the proposition." How does a sentence, which expresses a unified proposition, differ from a mere list of words, which does not? What does it take to convert the individual words into a unified, truth-evaluable representation? Burge (2005, pp. 18-21) argues that Frege's functional conception of semantic composition helps answer this question.

Specifically, Frege's conception of predication as functional application illuminates the propositional unity of atomic sentences. Davidson (2005, pp. 130-163) also commends Frege, although he thinks Tarski's more ontologically neutral approach provides greater illumination.

The unity of the proposition is a special case of a more general problem, arising for any representational system with recombinable parts (e.g. maps, diagrams, calendars, and music notation). How does a map differ from a mere list of symbols? What does it take to convert individual markers and coordinates into a unified, truth-evaluable representation? If my thesis is correct, we cannot answer this question by invoking the Frege-Tarski conception of predication. Thus, although the Frege-Tarski conception illuminates representational unity for sentences, it does not generalize into a satisfactory account of representational unity for maps.

How, then, should we explain the unity of cartographic representation? I submit that any decent answer should exploit the fundamental observation with which we began: a map purports to replicate relevant geometric features of physical space, including the spatial distribution of salient properties. In this respect, a map differs from a mere list of symbols. The symbols in a list may bear geometric relations to one another (e.g. one symbol lies a certain distance from another), but those relations typically lack representational significance.

A map's geometric structure is not just another element to be listed alongside its markers and coordinates. (Cf. Shin 1994, p. 161). That would engender a familiar regress: how does the map differ from a list containing markers, coordinates, and geometric structure? Rather, the markers and coordinates *stand in* geometric relations, and those relations bear representational import, as described by (3a-b). Markers and coordinates form a unified representation through being integrated into a geometric structure endowed with special representational significance.

I do not offer the foregoing analysis as a complete theory of cartographic representational unity. It raises many questions, such as what endows a map's geometric structure with representational import. Nor do I claim that the analysis is particularly novel. The basic idea stretches back to Wittgenstein's *Tractatus*. The key point is how much the analysis differs from the parallel analyses advanced by Burge and Davidson regarding linguistic representation. What binds a map into a unified representation of the world is the *geometric* structure that envelops its markers and coordinates, not the *function-argument* structure isolated by Frege and Tarski. Thus, cartographic and sentential representational unity derive from fundamentally different sources.[12]

A final possible application of my discussion concerns *mental representation*. Since at least the medieval era, commentators have often suggested that certain mental states represent in the same manner as pictures, diagrams, or maps. The modern literature most relevant for us discusses *spatial representation*. Beginning with Tolman (1948), many philosophers and psychologists have argued that animals navigate by employing *cognitive maps*: mental representations with the same basic representational properties and mechanisms as ordinary concrete maps. But what exactly does it mean for a mental representation to have "the same basic representational properties and mechanisms" as an ordinary concrete map? How would cognitive maps differ from sentential mental representations, of the kind posited by Fodor (1975) and many other researchers? A basic difficulty here is that concrete cartographic representation itself remains so poorly understood, especially in comparison with linguistic representation. By clarifying how cartographic compositional mechanisms differ from linguistic compositional mechanisms, my discussion may help clarify what cognitive maps are and how they differ from sentential mental representations. A natural suggestion is that cognitive maps are geometrically structured mental representations whose compositional semantics resembles that of ordinary concrete maps. If this paper's thesis is correct, then any such mental representation would differ profoundly in its representational format from sentential mental representations, whose semantics presumably resembles that of the predicate calculus. Developing these ideas is a task for another occasion (Rescorla forthcoming a), (Rescorla forthcoming b).

Notes

[1] Strictly speaking, we should distinguish between the *predicate* "*F*( )" and the *open-sentence* "*F*(*x*)" obtained by filling the predicate's open position with the variable "*x*." Although this difference is important for some purposes, it does not bear very directly upon my discussion, so I blur it for convenience.

[2] Many maps contain sentences, such as "You are here." A complete semantics for these maps will employ predicational mechanisms to analyze the relevant sentences. My position concerns the purely cartographic portion of these maps, i.e. what remains once one removes any sentences of natural or artificial language.

[3] Haugeland (1998, pp. 190-194) argues for a principled distinction between *logical* and *iconic* representations, the latter including pictures and maps. My discussion has some affinities with Haugeland's. However, Haugeland does not explicitly argue that maps lack function-argument predicational structure. He also places less emphasis on compositional mechanisms than I do.

[4] Maps have also received less philosophical attention than either pictures (Goodman 1976), (Hopkins 1995, 1998), (Kulvicki 2003), (Lopes 1996), (Malinas 1991), (Peacocke 1987), (Sober 1976), (Walton 1973) or diagrams (Allwein and Barwise 1996), (Hammer 1995), (Potter 2006), (Shin 1994).

[5] Citing projection mappings, Sloman (1985) argues that homomorphism should not play a prominent role in our theory of cartographic representation. However, I would urge that Sloman's argument relies on an overly restrictive conception of homomorphism. Virtually all cartographic projections preserve *some* relevant geometric structure. For instance, the Mercator and stereographic projections preserve the angles at which curves cross one another, while the Gall-Peters projection preserves relative areas.

[6] Note that the proposal does some violence to our intuitive understanding of markers. If we interpret a marker that denotes trees as a predicate of locations, then we must assign it a meaning like: "is occupied by trees." This strikes me as odd. Maybe we can live with such oddities.

[7] A related intuition runs as follows: if a marker appears at no map coordinate but is part of the mapping system (e.g. it appears in the map legend), then absence of the marker from a coordinate signifies absence of the corresponding property from the corresponding location. We may experience this intuition, but my arguments do not rely upon it.

[8] A third response holds that *alternative* conceptions of predication yield truth-conditions like those I ascribe to maps. I do not need to address this response here, since my thesis is simply that maps do not instantiate predication *as analyzed by Frege and Tarski*. The burden lies upon advocates of the response to delineate an alternative analysis of predication, to show how it handles the semantics of complex predicates and objectual quantifiers, and to show that it applies to maps. A fourth response insists that, even if my argument shows that attaching a marker to map coordinates is not predication, markers might nevertheless *be* predicates. Again, I do not need to answer this objection, since it concedes my main thesis. Moreover, it is difficult to see why we should classify markers as predicates once we concede that they do not participate in a predicative compositional mechanism. It seems clear that maps fall under nothing resembling the Tarskian clauses governing how the satisfaction relation interacts with the logical connectives. Thus, markers do not participate in *any* of the characteristic compositional mechanisms isolated by Frege and formalized by Tarski. In what sense, then, are they predicates? Lacking an answer to this question, the fourth response seems unconvincing.

[9] In cases such as M_2 where $\sigma_{1,C}$ is not topologically connected, deleting $\sigma_{2,C}$ yields a set of disconnected regions. It is not clear that we should call the result a "map." If we do, then the map is most naturally regarded as having a determinate truth-condition, a truth-condition that is neutral about whether $den(C)$ inhabits $den(\sigma_{2,C})$.

[10] Deleting "Mark" from "John shaves Mark" yields "John shaves", which has a truth-value. Such examples might seem to impugn my argument by calling into question the assumption that one can always delete a singular term from a sentence of the form "$\Psi(\alpha, \beta)$" to achieve an unsaturated formula "$\Psi(\alpha, \quad)$". In response, we might treat "John shaves" as containing an unarticulated reflective pronoun ("John shaves himself"). There is no clear reason to posit a comparable unarticulated constituent in our map semantics.

[11] We can imagine other putative counter-examples to the Absence Intuition, such as a relativized "*terra incognita*" symbol that signifies ignorance about a particular range of properties. It seems likely that, through maneuvers similar to those from §7.3.1-7.3.3, we can accommodate this and most other putative counter-examples.

[12] By juxtaposing these different sources, my account casts doubt upon theories, such as Sellars's (1979) Wittgensteinian account of predication, that analyze atomic sentences by analogizing or assimilating them to maps.

Acknowledgments

I am indebted to C. Anthony Anderson, Anthony Brueckner, Alex Byrne, Elisabeth Camp, Roberto Casati, Warren Goldfarb, Richard Heck, Paul Hovda, Luke Manning, Charles Parsons, James Pryor, and Nathan Salmon for helpful comments. I also received extremely useful feedback from audience members when I presented this paper at the Pacific APA 2006, and from my commentator on that occasion, William Taschek. Finally, I thank two anonymous referees for this journal for comments that helped improve the paper.

Works Cited

Allwein, G. & Barwise, J. (1996). *Logical Reasoning with Diagrams*. (Oxford: Oxford University Press).

Andrews. J. H. (1990). Map and language: a metaphor extended. *Cartographica*, *17*, 1-19.

Brachman, R. & Levesque, H. (2004). *Knowledge Representation and Reasoning*. (Boston: Elsevier).

Burge, T. (2005). *Truth, Thought, and Reason*. (Oxford: Oxford University Press).

Casati, R. & Varzi, A. (1999). *Parts and Places*. (Cambridge, MA: MIT Press).

Davidson, D. (2005). *Truth and Predication*. (Cambridge, MA: Harvard University Press).

Davis, E. (1986). *Representing and Acquiring Geographic Knowledge*. (London: Morgan Kaufmann).

Dilworth, J. (2002). Varieties of visual representation. *Canadian Journal of Philosophy*, *32*, 183-205.

Dummett, M. (1981). *The Interpretation of Frege's Philosophy*. (London: Duckworth).

Fodor, J. (1975). *The Language of Thought*. (New York: Thomas Y. Crowell).

Goodman, N. (1976). *Languages of Art*. (Indianapolis: Hackett).

Grice, P. (1989). *Studies in the Ways of Words*. (Cambridge, MA: Harvard University Press).

Hammer, E. (1995). *Logic and Visual Information*. (Palo Alton: CSLI Publications).

Haugeland, J. (1998). *Having Thought*. (Cambridge: Harvard University Press).

Hayes, P. (1985). Some problems and non-problems in representation theory. (In R. J. Brachman & H. J. Levesque (Eds.), *Readings in Knowledge Representation* (pp. 3-22). Los Altos: Morgan Kaufmann).

Head, C. G. (1984). The map as natural language: a paradigm for understanding. *Cartographica, 21*, 1-32.

Heim, I. & Kratzer, A. (1998). *Semantics in Generative Grammar*. (Malden: Blackwell).

Hopkins, R. (1995). Explaining depiction. *Philosophical Review, 104*, 425-455.

---. (1998). *Picture, Image, and Experience*. (Cambridge: Cambridge University Press).

Kitcher, P. & Varzi, A. (2000). Some pictures are worth $2^{\aleph_0}$ sentences. *Philosophy*, *75*, 377-381.

Kulvicki, J. (2003). Image structure. *Journal of Aesthetics and Art Criticism, 61*, 323-340.

Larson, R. & Segal, G. (1995). *Knowledge of Meaning*. (Cambridge, MA: MIT Press).

Lemmon, O. & Pratt, I. 1999. Putting channels on the map: a channel-theoretic semantics of maps? (In L. Moss, J. Ginzberg, & M. de Rijke (Eds.), *Logic, Language, and Computation, Vol. 2* (pp. 143-164). Stanford: CSLI Publications).

Leong, M.-K. (1994). *Towards a Semantics for a Visual Information System*. Dissertation, Stanford University.

Lopes, D. (1996). *Understanding Pictures*. (Oxford: Oxford University Press).

Malinas, G. (1991). A semantics for pictures. *Canadian Journal of Philosophy, 21*, 275-298.

Millikan, R. (2004). *Varieties of Meaning*. (Cambridge, MA: MIT Press).

O'Brien, G. & Opie, J. (2004). Notes towards a structuralist theory of mental representation. (In H. Clapin, P. Staines, & P. Slezak (Eds.), *Representation in Mind: New Approaches to Mental Representation* (pp. 1-20). Boston: Elsevier).

Peacocke, C. (1987). Depiction. *Philosophical Review, 96*, 383-410.

Potter, D. (2006). Diagrammatic representation in geometry. *dialectica, 60*, 369-382.

Pratt, I. (1993). Map semantics. (In A. Frank & I. Campari (Eds.), *Spatial Information Theory: A Theoretical Basis for GIS* (pp. 77-91). Berlin: Springer-Verlag).

Rescorla, M. (Forthcoming a). Chrysippus's dog as a case study in non-linguistic cognition. (In R. Lurz (Ed.), *The Philosophy of Animal Minds*. Cambridge: Cambridge University Press).

---. (Forthcoming b). Cognitive maps and the language of thought. *British Journal for the Philosophy of Science*.

Russell, B. (1923). Vagueness. (Rpt. in J. Slate (Ed.), *The Collected Papers of Bertrand Russell,* vol. 9 (pp. 147-154). London: Unwin Hyman).

Schlichtmann, H. (1979). Codes in map communication. *Canadian Cartographer, 16*, 81-97.

Sellars, W. (1979). *Naturalism and Ontology*. (Atascadero: Ridgeview).

Shin, S.-J. (1994). *The Logical Status of Diagrams*. (Cambridge: Cambridge University Press).

Sloman, A. (1978). *The Computer Revolution in Philosophy*. (Hassocks: The Harvester Press).

---. (1985). Afterthoughts on analogical representations. (In R. J. Brachman & H. J. Levesque (Eds.), *Readings in Knowledge Representation* (pp. 432-439). Los Altos: Morgan Kaufmann).

Sober, E. (1976). Mental representations. *Synthese*, *33*, 101-148.

Strawson, P. F. (1971). *Logico-Linguistic Papers*. (London: Methuen).

Tolman, E. (1948). Cognitive maps in rats and men. *Psychological Review*, *55*, 189-208.

Walton, K. (1973). Pictures and make-believe. *Philosophical Review*, *82*, 283-319.