Mind & Language ISSN 0268-1064
Vol. 9 No. 1 March 1994

Article

Lloyd's Dialectical Theory of Representation

KENNETH AIZAWA

Dan Lloyd's *Simple Minds* is an ambitious project ranging over numerous issues in cognitive science, from the neurobiology of Hermissenda and information processing in the retina, to the nature of mental representation, consciousness and human cognition. A central part of this project is the development of a naturalized semantics, a theory that explains how it is that events or objects in brains represent events or objects in the external environment. Lloyd's explanation is a version of correlational semantics he dubs the *dialectical theory of representation.* The cornerstone of the dialectical theory is the supposition that an event r represents the event o that has the highest conditional probability given r. Around this concept, Lloyd articulates three conditions that are best understood as a set of sufficiency conditions for r representing o.

The present paper is part of a series of papers on correlational semantics. A large part of this series (Adams and Aizawa, 1992, 1993, and forthcoming) develops objections to Jerry Fodor's latest version of correlational semantics, the asymmetric causal dependency theory (Fodor, 1987, 1990a). In these papers on Fodorian semantics, we object that various formulations of Fodor's purported sufficiency conditions on meaning are both too strong and too weak, not allowing items in human brains to mean what they do in virtue of Fodor's conditions and allowing meaningless objects to be credited with meaning. Further, it is argued that Fodor's theory suffers

The author would like to thank Fred Adams and especially Gary Fuller for comments on an earlier draft of this paper. The present paper was supported in part by the Robert Brown Memorial Readership at the Marine Biological Laboratory in Woods Hole, MA. This support is gratefully acknowledged. Address for correspondence: Department of Philosophy, Central Michigan University, Mt. Pleasant, MI 48859, USA. Email: 35c4ub2@cmuvm.bitnet.

from relatively familiar problems with making what are the intuitively correct content assignments to mental representations. The present paper raises several objections to Dan Lloyd's version of correlation semantics. Although Fodor's and Lloyd's theories are rather different, they suffer from many of the same problems. The conditions constituting Lloyd's theory are too weak and they lead to counterintuitive content assignments. A final paper in the series (Adams, Aizawa and Fuller, in preparation) takes a more positive line. It attempts to show how the latest version of Dretske's correlation semantics (Dretske, 1988) can solve the problems that have beset Fodor's and Lloyd's theories.

Lloyd's theory is worthy of consideration in this series of papers on correlation semantics, since it constitutes an elaborate development of a strategy for dealing with the content specificity problems that generally beset correlational approaches. It proposes that the incorporation of multiple input channels to a representation r will be an essential part of the solution to the proper assignment of content. This approach clearly diverges from Fodor's idea of relying upon the asymmetric dependence of one nomological connection on another. Moreover, it ignores the implicit appeals to explanatory history found in Dretske's 1988 teleo-informational approach. Its closest relative may be some ideas found in Dretske (1986). There Dretske mentions the possibility of multiple input channels for fixing content, but goes on to include a theory of learning. Lloyd rejects Dretske's appeal to learning and, in essence, attempts to use a multiple input channel condition, among others, to articulate an adequate theory of meaning.

The plan for the present paper is simple: after introducing the dialectical theory in Section 1, a separate section is dedicated to each of four problems facing the theory.

1. The Dialectical Theory of Natural Representation

For expository purposes, it is convenient to consider the dialectical theory alongside some of its closely related competitors: Dretske's information-theoretic semantics and Fodor's asymmetric causal dependency theory. Dretske (1981, 1988), at least in principle, bases his theory of the representational connection between an event or object in the head r and an event or object in the external environment o on an information connection: r carries information about o. Fodor (1987, 1990a), in contrast, bases his theory on a nomological connection between some property r and some property o. Lloyd, however, proposes to use *probabilistic dependencies* defined in terms of conditional probabilities. That is, rather than beginning with the condition that r carry information about o (Dretske) or stand in a nomological causal connection with o (Fodor), Lloyd suggests that we begin with the condition that the probability of event o given r be different than the probability of o alone. (In notation, $Pr(o/r) < > Pr(o)$.) So, for Lloyd, r represents o only if r probabilistically depends on o (plus some additional conditions to be specified); that is r represents o only if Pr(o/r) < > Pr(o) (plus some additional conditions to be specified). A few comments about these differences are in order. Where Dretske has the rather strong requirement that Pr(o/r) = 1, Lloyd proposes only that Pr(o/r) < > Pr(o). For Dretske, the requirement that Pr(o/r) = 1 may or may not capture an underlying lawful connection between o and r. Fodor, on the other hand, dispenses with all talk of probabilities and goes straight for the nomological causal connections between putative r's and putative o's. On this score, Lloyd sides with Dretske and abandons causal connections for probabilistic dependencies. For Lloyd, conditional probabilities are *definitive* of the dependencies involved in representation; they are not merely indicators of causal connections between r and o. Lloyd (1989, pp. 38–9) writes, 'It may turn out that the best analysis of causality will be in terms of dependency, as some philosophers have suggested. I propose to take a less tendentious route, and simply flag those probabilistic relations with their own terms.'

Having fixed upon conditional probabilities to define probabilistic dependencies, Lloyd faces an immediate problem. Conditional probabilities define a symmetric relation between events, but representation is an asymmetric relation. It is an elementary theorem of the probability calculus that Pr(a/b) < > Pr(a) if and only if Pr(b/a) < > Pr(b). This threatens to make a probabilistically dependent on b just in case b is probabilistically dependent on a. In contrast to probabilistic dependency, representation is an asymmetric relation: 'Fred' represents Fred, but Fred does not represent 'Fred'. Lloyd must, therefore, adapt conditional probabilities to make them serviceable for representation. To this end, he adds a condition that directs the dependence from one event to another. He proposes if Pr(a/b) > Pr(b/a), then b depends on a.

Yet another background condition governs the sort of event that occurs in the conditional probabilities involved in representation. The definition of a probability measure involves specifying certain simple events from which can be constructed more complex events through conjunction and disjunction. Thus, tossing a head with a quarter, tossing a tail with a quarter, tossing a head with a nickel, and tossing a tail with a nickel might be simple events. One can then speak of such complex events as tossing a head with a quarter and a tail with a nickel, or tossing a head with a quarter or a tail with a nickel. Lloyd, however, stipulates that the events that enter into the representation relation he defines must be real, *natural events*, rather than the *artificial events* that may be constructed by arbitrary conjunction and disjunction of simple events. Artificial events are necessary for defining a probability measure, but they do not, according to Lloyd, occur in representation.[1] He writes (ibid., p. 21):

[1] Lloyd does not explain that the restriction of events to natural events might be reconciled with the formal demands of defining a probability measure, but we may ignore this technical difficulty. Plenty of further difficulties will emerge.

> The basic mereology of a conceptual scheme . . . suggests a distinction between two classes of possible concatenations of events, partitioning the natural from the artificial events. Fertilization, gestation, and giving birth are parts of a composite event, reproduction. Giving birth and withdrawing $10,000 from a bank account is, prima facie, an artificial event. In the first case, we see the two component events as parts of a larger whole; in the latter, we do not.

And later (p. 63):

> A natural event is a whole composed of parts, and we refer to it as a single event. An artificial event, on the other hand, is a concatenation of two or more natural events where the two events are not themselves part of a more inclusive natural event; we refer to it either with a plural noun ('the earthquakes'), a conjunctive phrase ('the earthquake and attendant tidal wave'), or a special term coined just to cover the composite event.

Lloyd intends his conception of natural and artificial events to be an elaboration of, or further application of, the concept of natural and artificial kinds. Natural events are a (natural?) kind of natural kind. Artificial events are a (natural?) kind of artificial kind.[2]

The preceding constitutes the laying of foundations for the theory; it is thus comparable to Dretske's (1981) elaboration of a version of the mathematical theory of information. It constitutes an attempt to adapt conditional probabilities to define a relation of probabilistic dependence that will meet the needs of a theory of representation. As with all forms of correlational semantic theories, the bulk of the philosophical work comes in the next step, the addition of further conditions that convert the central dependency relation (e.g. informational, causal, nomological, probabilistic) into a meaning or representation relation. In other words, the tough step comes when correlational approaches must add conditions that promote their central dependency relation from what is, at best, a Gricean natural meaning relation to the level of a Gricean non-natural meaning relation. In still other words, the challenge is to move from the sense of 'means' in which a pattern of wear on a set of automobile tires means that the front end is out of alignment to the sense of 'means' in which an item in the mind, such as the firing of a neural circuit, can mean that the front end

[2] It is not clear whether Lloyd takes the requirement that events that enter into the representation relation be natural events to entail that only natural kinds can be represented on the dialectical theory. If so, then it would, with no apparent rationale, exclude the representation of artifacts, such as cars, personal computers and skyscrapers, from the scope of the theory. It would also exclude the representation of individuals, such as Ben Franklin and the state of Wyoming.

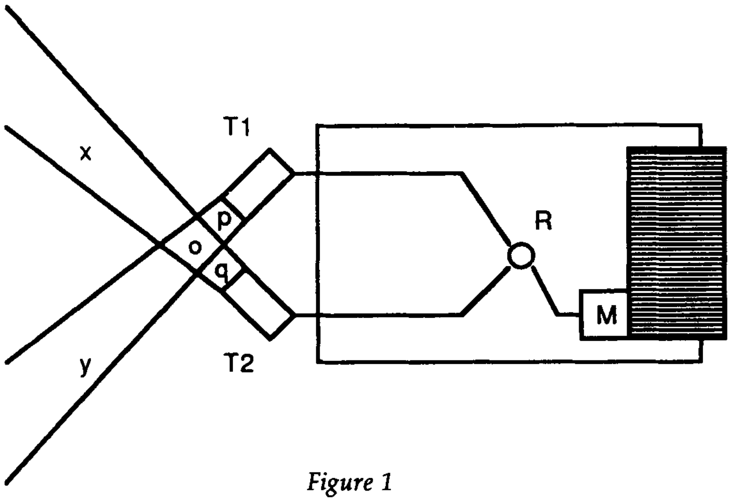

Figure 1

is out of alignment. Lloyd tries to effect this promotion by incorporating the probabilistic dependency relation into a set of three conditions. According to Lloyd, a natural event r is a representation if it meets the following conditions (1989, p. 64):

(1) *The multiple channel condition.* There is a set of at least two events, $\{v_1, v_2, \ldots, v_n\}$, such that r is dependent on the concurrent conjunction of at least two events in the set. (For example, r is a state change in an and-gate or a 'majority rule' threshold device . . . .)
(2) *The convergence condition.* Events v_1 through v_n are further subject to the constraint that there is a set of single events, $\{u_1, u_2, \ldots, u_n\}$ (the mutually effective stimuli), such that all of the v_1 through v_n depend on each element of $\{u_1, u_2, \ldots, u_n\}$. The object of a representation is the element of $\{u_1, u_2, \ldots, u_n\}$ with the highest conditional probability given r. (When conditional probabilities are tied, the representation is ambiguous.)
(3) *The uptake condition.* Event r has the capacity to cause either another representation or a salient behavioral event.

In part as an illustration of the theory and in part as a heuristic device to motivate the theory, Lloyd provides an example of a device that realizes his theory, the simple device called 'Squint' shown in Figure 1. Squint lives in a smooth, planar world and needs to move in the direction of light in order to survive. To this end, Squint has a structure R connected to a motor M in such a way that when R produces a signal r, Squint moves forward. R produces tokens of the representation r that are causally effective in bringing about behavior. This is how Squint satisfies the uptake condition.

Squint is supposed to fulfill the multiple channels condition in virtue of the fact that R is an and-gate.[3] R issues an r just in case it receives signals t_1 and t_2 from both the structures T_1 and T_2 at the same time. In the statement of the multiple channel condition above, Lloyd suggests that the condition requires that r be probabilistically dependent on the conjunction of t_1 and t_2 and that being an and-gate is merely one example of a way in which to have r be probabilistically dependent on t_1 and t_2. In other passages, however, Lloyd seems to hold that r's being probabilistically dependent on t_1 and t_2 is not sufficient and that R must be an and-gate. He writes (ibid., pp. 64-5; emphasis added):

> The multiple channel condition: This condition concerns the relation of a representation to 'upstream' or afferent events. It states that there are several upstream events, and that *the representation depends on two or more of them and will not occur otherwise.*

If Lloyd only wants to require that r be probabilistically dependent on t_1 and t_2 (i.e. $Pr(t_1 \& t_2/r) <> Pr(t_1 \& t_2)$), this does not entail that $Pr(t_1 \& t_2/\text{not-}r) = 0$. So, he has the and-gate interpretation in mind. The following passage reasserts this interpretation and gives an argument why this interpretation is necessary. Lloyd writes (ibid., p. 65; emphasis added):

> Not everything is an and-gate or majority rule threshold device. Although any event can be described as the consequence of a conjunction of events—with conjoint events sufficient for its occurrence—not every event has particular event conjunctions as a *necessary condition.*

The argument in this second passage is rather condensed. Here it is in a bit more detail. Requiring that r be probabilistically dependent on two events threatens to be vacuous. If r is probabilistically dependent on some event t_1, then it must be probabilistically dependent on some conjunction of events. So, let t^* be any event that is causally independent of r and t_1. If $Pr(t_1/r) <> Pr(t_1)$, then $Pr(t_1 \& t^*/r) <> Pr(t_1 \& t^*)$. To avoid this vacuity problem, it appears that in the passage just quoted Lloyd wants to abandon the idea that r be probabilistically dependent on t_1 and t_2 and adopt the requirement that the conjunction of t_1 and t^* be necessary and sufficient for r. This avoids the vacuity problem since, even if r is probabilistically dependent on the conjunction of t_1 and probabilistically irrelevant t^*, the conjunction of t_1 and t^* is not necessary for r. Not just

---

[3] For simplicity I will typically refrain from stating the multiple channels condition in its full generality allowing for more than two inputs to R.

any old t^* can be a necessary conjunct for r. Thus, rather than requiring that $Pr(t_1 \& t_2/r) <> Pr(t_1 \& t_2)$, the requirement of multiple channel condition appears to be r iff t_1 & t_2.

The structures T_1 and T_2 in Squint are light transducers and, although it is not strictly required by the conditions of the theory, it turns out that the transducers in Squint are not entirely reliable. While they usually fire only in the presence of light, they are also known to fire spontaneously in the absence of light. T_1 and T_2 are, however, causally independent, so that the T_1's firing does not influence the chances of T_2's firing, and vice versa. One consequence of this is that while each transducer has a small probability of misfiring, there is an even smaller probability that both will misfire simultaneously, thereby provoking R to issue an r token.

Consider, now, the convergence condition. This condition uses two ideas in an attempt to fix the object r will represent. The first idea begins with the observation that both T_1 and T_2 have visual fields. In Figure 1, O labels the overlap of the fields of T_1 and T_2, y labels the field of T_1 beyond O, x the field of T_2 beyond O, p labels the field of T_1 before O, and q labels the field of T_2 before O. Since T_1 is a light transducer, a light of sufficient intensity in the region p, o, or y is an *effective stimulus* for T_1. Similarly, a light of sufficient intensity in the region q, o, or x is an effective stimulus for T_2. In addition to effective stimuli, we might recognize *mutually effective stimuli,* those that will simultaneously trigger both T_1 and T_2. These include a light region in o, a light in region p conjoined with a light in region q, a light in region x conjoined with a light in region y, a light in region p conjoined with a light in region y, and a light in region q conjoined with a light in region x. Of these mutually effective stimuli, Lloyd wishes to count only a light in region O as a *single, mutually effective stimulus,* because, according to Lloyd, only a light in region O counts as a natural event. The conjunction of a light in region x and a light in region y, for example, is not supposed to be an object of representation in virtue of the conditions of the dialectical theory, since the conjunction is not a single natural event (a single mutually effective stimulus), but two events (ibid., p. 63). This is Lloyd's first means of fixing upon events in region O, rather than conjunctions of events in regions p, q, x, and y, as the events that are the objects of representations. His second idea is the requirement that the object of representation be the mutually effective stimulus with the highest conditional probability given r. To this he adds the empirical hypothesis that, in the relevant worlds, the mutually effective stimulus with the highest conditional probability is one that occurs in region O. In other words, given the set of mutually exclusive stimuli, {a light in region o, a pair of lights in regions x and y, a pair of lights in regions p and q}, not only is a light in region O the only single mutually effective stimulus, it is also the event that is most likely to have brought about r. Here is how Lloyd (ibid., pp. 69–70) makes the point:

> Earlier we noted that the multiple channel condition could be met by any of several mutually effective stimuli, of which o and (x and y) are examples. But note that even though o, x, and y are independently equiprobable, o is much more likely than the conjunction of x and y. I think the following generalization is plausible: By and large, single mutually effective stimuli events are more likely than complex (conjoint) mutual stimuli event. This is an empirical generalization, not a logical truth; but as we reflect on more complex cases (cases with more information channels and greater specificity), it seems more and more plausible.

The convergence condition, thus, has two ideas that seem to overdetermine events in region O being the objects of representation.

From the foregoing, one can see that the uptake condition concerns what happens from the structure R onwards, the multiple channel condition concerns the connection between the transducers and structure R, while the convergence condition governs events leading into the transducers. Although Lloyd uses Squint to illustrate his theory in a purely perceptual context, he intends to allow events in neural circuits embedded entirely within the brain to play the various roles required by his theory. In fact, this is the aspect of the theory that gives it its name. Within the brain there is a 'dialectical' alternation between convergence and divergence of signalling. As Lloyd puts it, 'Multiple channels, each with a slightly different "point of view", are synthesized in a single integrative device, in which changes represent single events. The process of synthesis can be repeated, with representing devices contributing their output to further downstream representing processes'.

Early in his book, Lloyd suggests that he will offer neither necessary nor sufficient conditions for representation. He claims that 'The traditional philosophical ideal of discovering logically necessary and sufficient conditions for the concepts of mind will not be appropriate here' (ibid., p. 11; cf. Lloyd, 1987, p. 24). Yet, when the philosophical work gets under way, he does turn to the traditional concepts of necessary and sufficient conditions. He sometimes speaks as if his three conditions constitute a definition of what it is for r to represent o, hence that they are necessary and sufficient for r to mean o. At other times, however, he suggests that the conditions are only sufficient for r representing o. Indeed, he can vacillate between these two interpretations quite rapidly (ibid., p. 62):

> I will first introduce a few technical distinctions, three special terms that will figure in the definition of representation. Then I will state the sufficient conditions for an event to be a representation. The hundred or so words of the definition will, not surprisingly, call for a lot of discussion.

Nonetheless, Lloyd's considered view appears to be that he has sufficiency conditions for r meaning o. Lloyd realizes that his theory cannot provide necessary conditions on meaning simpliciter, when he acknowledges that his theory cannot plausibly account for the meanings of abstract objects (ibid., p. 40). In addition, he seems to recognise that not all representations of physical objects get to be representations in virtue of the conditions of the dialectical theory. He asserts that 'one would . . . take a half-tone photograph of a small dot, producing a photograph in which one atomic dot represents another' (ibid., p. 84; cf. p. 167f.). A photograph is a representation that does not meet all of Lloyd's conditions. A photograph has no and-gate or 'majority rules' detector, and no structure R to produce tokens r. Nor do photographs have anything constituting the channels $\{v_1, v_2, \ldots, v_n\}$.[4] Finally, in Chapter 6 of *Simple Minds*, he claims that such things as negation and universal quantification cannot be captured with the dialectical theory.

Having explained what each of the conditions of the dialectical theory requires, and their collective status as a set of sufficient conditions for meaning, I should explain why Lloyd includes each of them in the dialectical theory. Consider, first, the multiple channel condition. Lloyd offers three reasons for including the condition. The first is based on principles of evolution by natural selection: by reflecting on a possible course of evolution, one could see how having multiple channels might be valuable (ibid., pp. 51ff., 60, 71).[5] The second is that, combined with the convergence condition, it is supposed to serve as part of a solution to the problems of specificity of meaning (ibid., p. 60–1). In fact, this second argument is intended to justify both the multiple channel condition and the convergence condition, since it is the combination of the two that solves the problems. The third justification is that, in combination with the convergence condition, the multiple channel condition provides for a theory of error.

Consider, first, the evolutionary justification. In the environment in which Squint's ancestors found themselves, it was important that they move forward in the presence of light, but not otherwise. These ancestors had only one light transducing channel with a structure R that served merely as a relay node, rather than an and-gate. These ancestors are illustrated in Figure 2. The transducers and channels in Squint's ancestors, like those in Squint, were unreliable because they were prone to misfire spontaneously, that is, they would signal the presence of light when in

[4] Lloyd (1987, pp. 27, 50) treats an earlier version of the conditions as both necessary and sufficient for representation, despite the fact that this photograph counterexample threatens to serve as a counterexample there as well.

[5] The evolutionary story actually does double duty for Lloyd. In the initial stages of exposition it serves merely as a heuristic device for introducing the condition. In the later stages cited above, it is repeated in order to justify the inclusion of the multiple channels condition in the dialectical theory.

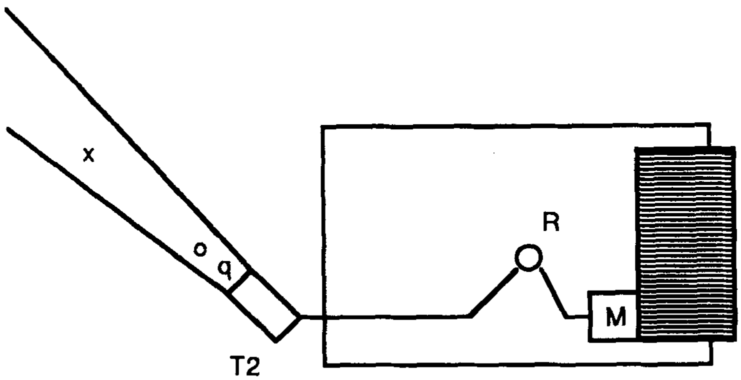

Figure 2

fact no light was present. In response to such misfires Squint's ancestors might have evolved a more reliable single channel, one with a smaller probability of spontaneous misfirings, but they did not. Instead, Squint's ancestors evolved two transducers, each with its own channel that feeds into the sort of and-gate found in the contemporary Squint. This two-channel arrangement increases the reliability of the light indicator R just as well as can improving the single channel, since the probability of both channels spontaneously misfiring at the same time is less than the probability of a single channel spontaneously misfiring. Once this two-channel system appeared in the gene pool, the increased reliability of light detection and consequent increase in fitness quickly drove the trait to fixation in the population. Thus, Squint himself possesses this two-channel system rather than the equally workable system having an improved single channel.

In Chapter 2 of *Simple Minds*, in his criticism of Dretske's approach to representation, Lloyd observes that one of the central problems for the correlational approach to semantics is to specify which event constitutes the object r represents. Of all the natural events with which r is correlated, which does r represent? A more specific version of the question asks which natural event in a single dependency chain leading from r into the indefinite past does r represent. Consider an example using Squint. Given that r represents what it is probabilistically dependent upon, why is that r represents a light in region O, rather than, say, the photons in region O at a distance of 1 micron from the surface of the light, the photons in the path between the region O and the two transducers, the flow of current in the circuit of which the light bulb is a part, the closure of an electrical circuit, a person's flipping a switch, or a person's intention to flip a switch? Lloyd refers to this as the *linear specificity problem*. Dennis Stampe (1979) called it the *diachronic specificity problem*.

The second justification for the multiple channel condition is that it is part of a solution to this problem. Lloyd (1989, pp. 60–1) claims that:

> Squint's crossed photodetectors provide one plausible model for mechanisms with linear selectivity. A single photocell cannot distinguish between a light source two inches away and a much brighter source at two feet, so it cannot carry information about the distance or the actual brightness of any source. But a device such as I've described not only serves to help confirm a signal [hence, increasing reliability], it is selectively sensitive to specific distal event types. When [R] is picking up information from both of its two input channels, then the upstream dependent event with the highest probability is the conjoint cause of both inputs, a light source at o. A multiple channel mechanism, in effect, focuses on its object.

He later continues (ibid., pp. 70–1)

> Suppose we delete the multiple channel condition. Figure [2, above] illustrates the results of a channelectomy. (We assume that r is no longer instantiated in an and-gate or majority rule threshold device.) The price of this deletion is high. In this case, the occurrence of r tells us nothing about the likelihood of any specific upstream event; it is equally an indicator of t_2, o, and x. We thus lose any special purchase on accuracy or selectivity, reintroducing the problem of linear selectivity.

Lloyd's solution to the linear selectivity problem involves the multiple channels condition, since one evidently needs multiple channels to define the mutually effective stimuli. The convergence condition enters through the appeal, in the first passage, to the event leading up to t_1 and t_2 having the highest conditional probability. Thus, according to Lloyd, r represents a light in region O, rather than, say, the photons in the path between the region O and the two transducers, the flow of current in the circuit of which the light is a part, the closure of an electrical circuit, a person's flipping a switch, or a person's intention to flip a switch, because a light in region O is supposed, as a matter of empirical fact about the environment in which Squint finds himself, to be the single, mutually effective stimulus with the highest conditional probability given r.

The linear specificity problem is only part of the challenge of saying which natural event of those with which r is correlated r represents. There is in addition to the linear problem the *lateral problem*, what Stampe called the *synchronic specificity problem*. Consider some event e in the causal chain leading up to the production of r. Which of all the events occurring at the same time as e counts as the event represented? For example, suppose a light bulb is on in region O of Squint's visual field. At the same time the light shines, a massive object generates a gravitational field, a filament oxidizes, heat is radiated, entropy is increased, and so forth. Which of these events does r represent? To answer the question, Lloyd observes that 'Only light can provoke T_1 and T_2, and thus [R]. The device is insensitive to many features of the stimulus (its temperature, for example)' (ibid., p. 61). In other words, the firing of r represents the light being on, since the transducers are selectively sensitive to light, not gravity, oxidation, heat, or the increase of entropy. In other words, it is qua light generator that the bulb sets off the detectors, hence r must represent the light. It should be noted that this answer to the lateral selectivity problem does not involve the multiple channel condition. The property selectivity of the transducers T_1 and T_2 could be had as easily with one transducer as with two or more. Evidently, if the property selectivity of the transducers is to be tied to any of the conditions of the dialectical theory, it must be to the probabilistic dependencies of the convergence condition.

This brings us to the third justification for the multiple channel condition, an argument that in fact supports both the multiple channel condition and the convergence condition. In combination, the conditions allow for errors of commission, where an organism thinks something that is false. The challenge is the following (cf. Dretske, 1981, ch. 8; Fodor, 1990b). How can it be that, on a correlational theory, event r is probabilistically dependent on o', without o' being a part of the meaning of r? Suppose Fred is a normal adult male driving down the expressway. He glances out of the window of his car, sees some animal in the field, and thinks, 'cow', but the animal he glimpses is not really a cow, but a horse. Here one intuitively wishes to say that Fred was in error; he mistook a horse for a cow. If, however, syntactic tokens of 'cow' mean what they do in virtue of what they are probabilistically dependent upon, then what principled reason is there for saying that 'cow' means cow, rather than cow or horse? In 1981, Dretske answered this sort of challenge in terms of learning conditions, and in 1988, in terms of what 'cow' has acquired a function to indicate. In 1987 and 1990, Fodor appealed to a purported asymmetric causal dependency of a horse-'cow' law on the cow-'cow' law. Lloyd's theory allows for errors of commission as follows. In organisms such as Squint, a conjunction of events in the regions x and y will set off both transducers T_1 and T_2, hence trigger r, but these count as errors. They are errors because r must represent in a single, natural event, such as 'There is a light in region O', but when there are lights in regions x and y, there is presumably nothing in region O. r, therefore, says something false.[6] Note that in this solution, Lloyd cannot claim that 'cow'

[6] Lloyd also has a theory of errors of omission, not saying that there is a light in region O, when there in fact is. This sort of thing might happen when Squint is fog-bound so that no light from the region O reaches Squint's transducers. This theory of errors of omission does not help justify the multiple channel condition or the convergence condition, since it is independent of both. Even a single transducer can suffer errors of omission of this sort.

means cow, because it is more likely that a cow is present given r than it is that a cow or horse is present given r. This fails for the simple technical reasons that $Pr(t_1/r)$ is always less than or equal to $Pr(t_1 \text{ v } t_2/r)$.

I have just shown the justifications Lloyd offers for the multiple channel condition and the convergence condition. This leaves me with the task of explaining the rationale for the uptake condition, that the representation r must have the capacity to cause either another representation or a salient behavioral event. Squint illustrates the third condition insofar as r events in the structure R will cause Squint to move forward. Lloyd's justification for the condition is not that behavioral effects are necessary for an object to be a representation *per se*. It is merely that without the condition, the representation would be explanatorily idle (1989, pp. 71–2):

> Finally, . . . we block the effects of r, rendering it inert. Two crucial conditions on representation are met, but what's the point? Event r is a dead end, so it cannot figure in an explanation of behavior, either directly . . . or through its influence on other representations . . . . Event r would be a curious one, but not a factor in any explanation of behavior and so not a strong candidate for a special role in psychology.

In this regard, the dialectical theory differs from Dretske's latest offering (Dretske, 1988). On Dretske's account, having behavioral effects is essential to an event's becoming representational. An event that indicates some environmental state of affairs could not become a genuine non-natural representation at all unless it had behavioral effects.

As a final bit of exposition, I might add just two very brief comments on the relationship between the dialectical theory and Dretske's earlier account (Dretske, 1981). In the first place, Dretske's account did not require multiple input channels for representation, so that some events that Dretske might count as representations will not be so counted by the dialectical theory. Second, Dretske believed (and continues to believe) that there are no innate representations, no innate concepts. Nothing in the dialectical theory, however, prevents there from being innate representations. The dialectical theory is, therefore, ahistorical in a way that Dretske's 1981 theory is not. So, even though Dretske's 1981 theory may have inspired the dialectical theory, the dialectical theory is certainly at odds with its progenitor.

The following four sections describe four objections to the dialectical theory of representation. First, Section 2 argues that Lloyd's conditions are not sufficient for representation. The conditions count some meaningless events as meaningful. Section 3 shows that reliability, Darwinian fitness, and the multiple channel condition do not go hand in hand in the way Lloyd suggests through his appeal to his selectionist story. Section 4 provides reasons to believe that the dialectical theory's treatment of the

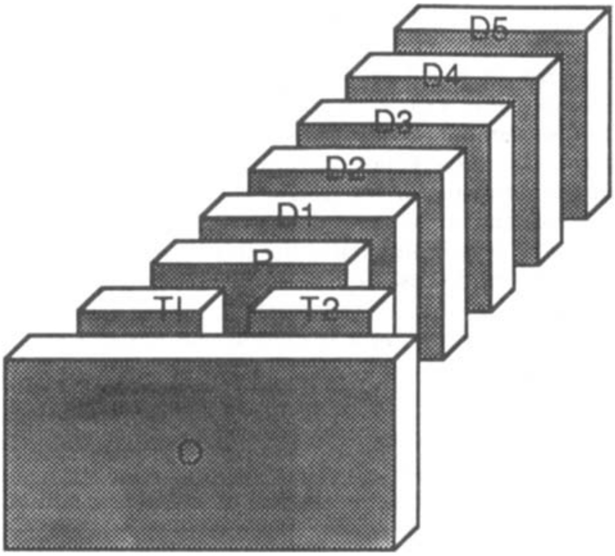

Figure 3

problems of specificity is ultimately inadequate. Finally, Section 5 shows that the dialectical theory of error is inadequate.

2. *Lloyd's Conditions are not Sufficient for Representation*

The most obvious sort of objection to a set of sufficiency conditions displays some system that meets the conditions, but which does not have the property for which the conditions are purported to be sufficient. So, I will begin with this sort of objection. Here's a system that satisfies the conditions of the dialectical theory, but contains no meanings. Consider the arrangement of dominoes shown in Figure 3. Let o be the event of domino O falling. When, and only when, O falls both dominoes T_1 and T_2 fall. Thus, events t_1 and t_2 are probabilistically dependent on event o. Next, domino R falls when, and only when, both T_1 and T_2 fall. R is thus an and-gate. Finally, when and only when R falls, the string of subsequent dominoes D_1, D_2, . . . D_5 falls as well. Perhaps this set of dominoes can be placed in a motorized box in such a way that when D_5 falls it starts the motor moving the box in the direction that D_5 fell. This system satisfies the three conditions of the dialectical theory, but r does not semantically mean that O fell. The event r naturally means that O fell, but it does not reach the goal of giving us non-natural meaning. The event r carries the information that O fell, but it does not mean that O fell. Put this in yet another way. The three conditions of the dialectical theory do not make r mean o in a more interesting sense than that in which domino D_2's falling means that domino D_1 fell. Both seem to be cases of mere information transmission, purely natural meaning. So, there seem to be meaningless events that the dialectical theory counts as meaningful.

From the preceding it should be clear that something is wrong with the dialectical theory, but a counterexample is merely the beginning. Saving the theory requires more than a little fiddling with the conditions here and there. The theory faces more fundamental difficulties. The remainder of this paper tries to bring forth these more fundamental difficulties.

3. Reliability and the Multiple Channel Condition

Lloyd takes increased reliability, hence an increase in Darwinian fitness, to provide some justification for the multiple channel condition. The selectionist story serves a heuristic function, but it also serves a justificatory function. Here we find that, rather than providing some justification for multiple input channels, there is in fact a tension between reliability of object detection and the possession of multiple channels.

Lloyd takes the possession of multiple channels to be important to representation, but he takes this importance to lie in the reliability of representation multiple channels confer. Another way to put this is to say that, according to Lloyd, having multiple channels is not intrinsically good for representation, not good in itself. Having multiple channels is good as a means to an end: it is good insofar as it increases the reliability of detection of events or objects that are relevant to the survival of the organism. This is clear from the way in which Lloyd introduces multiple channels in his conceptual evolution of Squint (ibid., p. 53f). It is also apparent in his explanation of why the loss of a channel is bad for a toad (ibid., p. 71). Having only one channel is not bad in itself; it is bad because it makes the toad's chances of detecting food somewhat lower.

The tension arises because there is no necessary connection between reliability and multiplicity of channels. A thought experiment should make this clear. Suppose Squint is in an accident and is reduced to the state of one of his one-eyed ancestors shown in Figure 2 above. Lloyd would (I will assume) claim that the 'channelectomized' Squint can no longer semantically represent events in its environment. Squint merely has an information channel. Suppose we can perform at most one of two possible surgical operations in an attempt to restore Squint's representational capabilities. In the first operation, we can divide Squint's remaining channel into two parts and use the parts to construct two channels. Suppose this is the only way to give Squint a new channel. This will give Squint two channels, but unfortunately, dividing a single channel in two makes it less reliable, more subject to noise, so that r spontaneously misfires more often with the two channels than it would have if we had left the single channel alone. Have we succeeded in restoring Squint's representational capabilities? If we look at Lloyd's three conditions in isolation from the reasoning that moved him to them, we would have to say yes. Squint now has multiple channels with a set of single mutually exclusive stimuli that contribute to the production of behavior. If, on the other hand, we look at the reasoning underlying the conditions, the consideration of reliability, we would have to say no. The operation renders Squint less reliable. The standards of reliability and multiple channels do not agree in their assessments of Squint's representational capacities.

Consider the second possible operation. Rather than damage Squint's one good channel in making two new channels, we can improve it. We can make it more reliable, so that it will spontaneously misfire less often. Unfortunately, taking the measures needed to produce the single superior channel precludes any possibility of giving Squint a second channel. If we can make these repairs, will we restore Squint's representational capabilities? On the one hand, he still has only one channel, so according to condition (1), Squint fails to meet a necessary part of Lloyd's sufficient conditions for representation, so the operation is a failure. But, he is more reliable, so perhaps he does have some representational capabilities? Again, considerations of reliability and multiple channels disagree in their assessments. No matter which operation we contemplate, we find a conflict between the requirements of multiple channels and reliability.

The preceding thought experiment constitutes one way of seeing the conflict between the multiple channels condition and reliability. Here is another. It is a standard metatheoretic principle to suppose that when one offers a set of sufficiency conditions, each of the conditions is necessary for the sufficiency conditions. There should be no gratuitous conditions in a set of sufficiency conditions; removal of one condition should destroy the sufficiency of the conditions. This is so, despite the fact that, technically, a set of sufficiency conditions with extra gratuitous conditions can still constitute sufficiency conditions. To put the conflict between reliability and multiple channels another way, then, I note that, contrary to Lloyd's suggestion (ibid., p. 71), against the backdrop of the other sufficiency conditions, multiple channels are not necessary for reliability of object detection even in the organism Squint. As I noted in the exposition of Lloyd's view, and as Lloyd himself noted (p. 54), it is possible to increase the reliability of object detection in Squint simply by building a better single channel. In other words, the multiple channel condition is not a necessary part of a set of sufficient conditions on representation.

Faced with a conflict within his theory, Lloyd might offer a number of responses, but I wish to focus on only two of the most obvious possibilities: (1) reject the view that multiple channels are a necessary part of the set of sufficient conditions on meaning, but maintain the view that reliability is part of the set of necessary conditions on representation; or (2) reject the importance of reliability to representation, but maintain the importance of multiple channels.

Could Lloyd embrace the first alternative? Is the existence of a reliable connection between r and o (leading ultimately to increased survivability) essential to r representing o? Here is a problem with supposing that it is. Representation does not typically come in degrees; things are not more or less representational. Reliability, on the other hand, does come in degrees. Cars vary in their degree of reliability. Forms of birth control vary in their degree of reliability. We therefore need some principled means of saying that some degree of reliability is sufficient for representation, where some lesser degree is not. Lloyd cannot say that representation requires that the connection between r and o must be perfectly reliable, that is, that Pr(o/r) = 1, since he explicitly rejected this in his discussion of Dretske (ibid., pp. 42 ff). He must, therefore, try to muster a principled distinction between degrees of reliability between 0 and 1. He must say why, for example, 67%, or 75%, or 90% reliability is high enough for representation, but some lower value is not. This is not very promising.

Could Lloyd abandon talk of reliability and simply propose that the use of multiple channels is a necessary part of sufficient conditions for representation? Could he suggest that it is merely a brute fact that multiple channels must be a part of his sufficiency conditions? It would seem not. Consider a person with hearing in only one ear. Such a person can describe various features of a trumpet solo, carry on a conversation, and so forth, in much the same way as a normal adult. Voices and trumpets playing are effective stimuli for the ear drum, the nerve from the ear to the brain constitutes a single channel to various structures in the brain that form mental representations of the playing of musical instruments, the voicing of parts of speech, and so forth. The descriptions of the trumpet solo or responses to a query by a speaker constitute behavior made in response to the mental representations. It thus appears that a person with hearing in only one ear satisfies both the convergence and uptake conditions, although not the multiple channel condition.

One might observe that many individual neurons are involved in the interpretation of the vibrations sounds make on the ear drum and bones of the ear and think that these constitute the multiple channels required by the multiple channel condition. The problem with this response is that, while it yields representations that meet the conditions of the dialectical theory, the conditions dictate the wrong objects of mental representation. The neurons that are the multiple input channels do not focus on events in the external world, such as voices and trumpet solos. Instead, they focus on vibrations in the ear drum, or perhaps the bones of the ear. A person with one deaf ear would not, then, have a representation of a friend's greeting, 'Hi, Ken, how are you doing?', but a representation of a vibration in the ear drum. A person's idea of the trumpet playing would really be an idea of the properties of vibrations in the ear. This response seems to assign all the wrong objects of representation.

The preceding considerations would seem to constitute serious challenges to analyzing representation in terms of reliability or some condition on multiple input channels.

4. *The Specificity Problems*

Representations pick out specific parts of reality. They focus on individual objects, natural kinds, artificial kinds, and properties. A token of 'Fred' typically refers to Fred and no one else. A token of 'red' typically refers to, or picks out, the property of being red, rather than some individual person, a kind of animal, or the property of being green. Naturalistic semantic theories must capture this characteristic of representation. Getting these matters right was discussed above in terms of the linear and lateral specificity problems. Lloyd (ibid., pp. 44f) tested Dretske's 1981 theory on its ability to handle these problems, and concluded that while Dretske's account was at least somewhat successful in handling the lateral problem, it could not adequately solve the linear problem. It is, therefore, appropriate that we similarly test Lloyd's theory.

Let us return, first, to the problem of linear specificity. Here one must explain why r in Squint represents a light in region O, rather than, say, the photons in the path between the region O and the two transducers, the flow of current in the circuit of which the light is a part, the closure of an electrical circuit, a person's flipping a switch, or a person's intention to flip a switch. Lloyd cited two reasons for saying that a light in region O was the object of representation. First, he asserted that the object of representation must be a single event, such as a light in region O, rather than a conjunction of distinct events, such as a light in region x and a light in region y. Second, he claimed that the object of a representation must be the event with the highest conditional probability given r, and that, as a matter of empirical fact, a light in O is this event. A light in O is more probable than, say, a pair of lights at x and y, a pair of lights at p and q, the flow of current in the circuit of which the light is a part, the closure of an electrical circuit, a person's flipping a switch, or a person's intention to flip a switch.

Despite the double-barrelled defense of lights in O as the objects of representation, the account will not work. Consider what was, from the beginning, a gratuitous empirical assumption on Lloyd's part (pp. 69–70), namely, that the distal event of a light shining in region O is more probable given r than is the more proximal event of the light from region O to the transducers given r. This assumption is in fact false. The probability of light passing from O to T_1 and T_2 given r must be greater than the probability of a light in region o given r. If light from O is to trigger r, then it must pass from region O to the transducers. Thus, without the false empirical assumption, it does no good to claim that the object of representation must be the event with the highest conditional probability given r. It also does no good to claim that light passing from O to T_1 and T_2 is not a single natural event. Perhaps a light shining in region p and a light shining in region q count as two distinct events, but light passing from O to T_1 and T_2 *is* a single natural event by the very standards Lloyd gives us. Note that the phrase, 'light passing from O to T_1 and T_2' is not

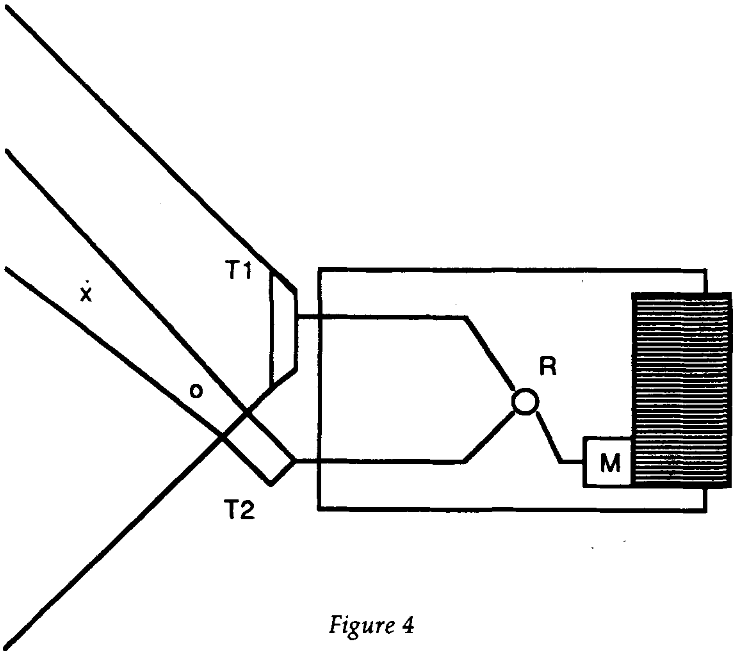

Figure 4

a plural noun (e.g. 'the earthquakes'), a conjunctive phrase ('the earthquake and attendant tidal wave'), or a special term coined just to cover the composite event, so it does not have the marks of an artificial event.

There is, as well, a more subtle problem with Lloyd's trying to solve the linear selectivity problem with multiple channels focusing on region O. Recall what Lloyd says of Squint should he lose his T_1 channel as shown in Figure 2. He says (ibid., p. 71), 'The price of this deletion is high. In this case, the occurrence of r tells us nothing about the likelihood of any specific upstream event; it is equally an indicator of t_2, o, and x. We thus lose any special purchase on accuracy or selectivity, reintroducing the problem of linear selectivity'. Not only does this version of Squint fail to solve the linear selectivity problem, it in fact has no representational system at all. It has only an informational system. Now, if the one-eyed version of Squint shown in Figure 2 cannot solve the linear selectivity problem and cannot be representational, then how can a two-eyed version shown in Figure 4? The version of Squint in Figure 2 and the version in Figure 4 have the same visual field. In the version of Squint shown in Figure 4 'the occurrence of r tells us nothing about the likelihood of any specific upstream event; it is equally an indicator of t_2, o, and x'.[7] The

[7] A version of Squint where the visual field of a first transducer is exactly the same as the visual field of a second transducer might make the point equally well. The only challenge appears to arise in fiddling with technical details to make two distinct transducers have the same conjoint field as each has individually.

problem is that nothing in the three conditions of the dialectical theory requires that a device with multiple channels have a narrower sensory field than a device with a single channel, and barring such an explicit stipulation, there is no necessary connection between adding input channels and decreasing the size of sensory fields. By Lloyd's reckoning, the dialectical theory admits what would seem to be metatheoretically excluded, namely, systems of representation that do not have linear selectivity. To put this another way, there is another system (recall the domino case) that meets the conditions of the dialectical theory of meaning, but which by Lloyd's accounting does not have meaning. Thus, there are two independent reasons for believing that Lloyd's second justification for the multiple channel condition does not work.

Let us turn now to the dialectical theory's handling of the lateral specificity problem. Semantic theories, such as Fodor's, that begin with the concept of nomological connection begin with a concept that displays a form of selectivity that Lloyd used in an attempt to explain lateral specificity. Nomological connections hold between specific properties. Suppose Fred has a compass C whose needle is aligned with magnetic north and that he then brings a magnet M near C and the needle changes its orientation. What caused the needle to move as it did? M's magnetic properties and not its mass, color, or density. The law covering the behavior of C in the presence of M was property specific. It relates M's magnetic properties to the metallic properties of C's needle. Nomological causal connections have in them this sort of property selectivity. It is part of their nature. This is some basis on which to build the lateral selectivity of representations. Lloyd apparently wishes to use the same sort of idea to solve the lateral specificity problem. He claims that a token of r in Squint represents light, because the transducers T_1 and T_2 in Squint were sensitive to light, rather than temperature, mass, and so forth. In other words, the law connecting r to objects in the external environment covered the property of being a light, and no other property of the events leading to tokens of r.

A problem begins to emerge, however, when we recall that the dialectical theory is based on conditional probabilities, rather than covering laws. Conditional probabilities do not naturally, or intrinsically, have the sort of property selectivity that natural laws do. The metaphysics of conditional probabilities is different from the metaphysics of covering laws. There is clearly nothing false or conceptually problematic in any of the following:

(a) Pr(C's needle moved/a magnetic object is brought near C) > Pr(C's needle moved);
(b) Pr(a magnetic object is brought near C/C's needle moved) > Pr(a magnetic object is brought near C);
(c) Pr(C's needle moved/ a massive object is brought near C) > Pr(C's needle moved);

(d) Pr(a massive object is brought near C/ C's needle moved) > Pr(a massive object is brought near C).

Yet, it would be false or misleading to say that C's needle moved because a massive object was brought near it. Compass needles don't respond to massive objects, they respond to magnetic objects. In truth, C's needle moved because a magnetic object was brought near it. The point, then, is that Lloyd cannot use conditional probabilities as they stand to introduce property selectivity.

It might be thought that the maximal conditional probability clause of the convergence condition would be of some help here. Perhaps we can say that, since Pr (a magnetic object is brought near C/ C's needle moved) is greater than Pr (a massive object is brought near C/ C's needle moved), the movement of C's needle probabilistically depends on the magnetism of M, not its mass.

This approach may at first appear to be promising, but it is in fact highly problematic. The problem lies in developing a consistent philosophical interpretation of the highest conditional probabilities. How are we to understand which conditional probabilities count as the highest? This is not an empirical question about the conditional probabilities of events in the world, but a conceptual question as to which probabilities to look for. Two sorts of cases suggest that an adequate interpretation will be hard to come by. First, why, on Lloyd's theory, is it not the case that r's only denote some of the most general properties of material objects, such as existing in spacetime, being massive, or being physical? Will it not be the case that Pr (an object in spacetime is present/r) > Pr (a bird is present/ r), Pr (a massive object is present/r) > Pr (a bird is present/r), Pr (a physical object is present/r) > Pr (a bird is present/r), and so forth? So, will it not be the case that nothing could form a concept of, say, a bird in virtue of satisfying the conditions of the dialectical theory? Why is it not the case that, on Lloyd's theory, all our representations are of the most general features of reality?

Consider the second, more involved case based on a modification of Squint. Let Squint's light detectors be sensitive to cardinals. But, further, suppose that Squint's transducers are such that Pr (a male cardinal is present/r) > Pr (a female cardinal is present/r) and that this is because males are more brightly colored than are females, hence are more easily spotted. In fact, we might suppose that Pr (a large male cardinal is present/ r) > Pr (a small male cardinal is present/r), since, other things being equal, the larger male cardinal will subtend a larger portion of Squint's visual field and will therefore be more likely to set off his transducers T_1 and T_2. Furthermore, we might suppose that Pr (a large male cardinal is two feet away/r) > Pr (a large male cardinal is ten feet away/r). This again will be because, *ceteris paribus*, the closer cardinal will subtend a larger portion of Squint's visual field. The natural consequence of this progression is that a system represents only the most perceptually salient events in its environment. Thus, tokens of r caused by suboptimal events, e.g. small male cardinals or female cardinals, will be erroneous, even when the production of r initiates what would seem to be appropriate behavior. If an occurrence of r is caused by a female cardinal and initiates Squint's predatory behavior (and, let us say Squint can and does eat female cardinals and gains nourishment from them), then that occurrence of r is still false, since r means 'large male cardinal two feet away' and no large male cardinal is around. Thus, few beings will have the concept of a cardinal in virtue of satisfying the conditions of the dialectical theory.[8]

So, the challenge is this. To save Lloyd's attempt to solve the lateral selectivity problem by appeal to second clause of the convergence condition, one must give a conceptual account of what is required by having r represent the event in the environment with the highest conditional probability that does not have a consequence either that only the most general or the most salient features of reality are represented. Note that what is needed here as a response is a single interpretation of the conditional probabilities that gets around the case of very general properties and the case of salient properties, while still handling the original problem of linear specificity. Using one interpretation to avoid one case and another interpretation to avoid another is insufficient. One interpretation must fit all. The upshot of this section is that Lloyd does not yet have available the means for explaining either linear or lateral representational selectivity.

5. The Theory of Error

In my exposition of the dialectical theory, I introduced three justifications for the multiple channel condition. First, it was purported to increase the reliability of object detection in systems prone to spontaneous firings. Second, it was part of a solution to the linear selectivity problem. Third, it was intended to help account for the possibility of error. According to this theory, errors of commission occur when r fires in response to the

[8] One might think that there is a third problematic interpretation of the highest conditional probability rule. Clearly, Pr (large male cardinal two feet away or Q/r) > Pr (large male cardinal two feet away/r) for arbitrary Q. If Pr (large male cardinal two feet away or Q/r) > Pr (large male cardinal two feet away/r), then r means large male cardinal two feet away or Q. If Pr (large male cardinal two feet away or Q/r) = Pr (large male cardinal two feet away/r), then r is ambiguous between 'large male cardinal two feet away or Q and large male cardinal two feet away'. In either case, the result is unsatisfactory.

Lloyd, however, avoids this problem through the requirement that the objects of representation be natural events, that is, the logical conjunction or disjunction of two natural events does not necessarily produce a natural event. The disjunction of there being a large male cardinal two feet away or it raining on Tuesday does not constitute a natural event, hence that we can exclude the event of raining on Tuesday from the extension of r. In such cases, the meaning remains large male cardinal two feet away.

combination of events in regions x and y, rather than in response to an event in region O. I have argued for the inadequacy of these first two justifications, so I might finally inquire briefly about this third.

According to the dialectical theory, the object of representation is supposed to be the event with the highest conditional probability given r. Lloyd then adds the empirical hypothesis that this event is some single mutually effective stimulus in region O. In my discussion of the linear selectivity problem, however, I noted that proximal single, mutually effective stimuli were more probable given r than were distal, single mutually effective stimuli in region O. The result of this, of course, is that not only are pairs of events in regions x and y counted as errors, but so are events in region O. This is unsatisfactory. To see a second problem, recall that by Lloyd's account, the one-eyed Squint of Figure 2 cannot distinguish between events in regions p, o, or x, hence it does not have the linear selectivity that Lloyd uses to introduce error. Once again, suppose that we add a second channel to Squint in the way shown in Figure 4 above, so that the region of overlap of the two transducers is simply the visual field of the single transducer T_1. This version of Squint satisfies the three (modified conditions) of the dialectical theory, but does not increase the selectivity of Squint's visual field and does not allow for the sort of error Lloyd envisaged on the dialectical theory. Thus, insofar as the ability to make errors of commission are necessary for a system to be a representational system, we have a case in which a system meets the sufficiency conditions, but is not a representational system. The preceding arguments, of course, parallel the arguments brought forth against the dialectical theory's handling of the linear selectivity problem, so without a solution to the linear selectivity problem, there is no solution to the problem of error.

6. Conclusions

A number of arguments have been brought to bear against the dialectical theory of representation. Most obviously the conditions fail to be sufficient for meaning. The domino counterexample shows this. Other objections were also brought forward in an attempt to indicate that more is needed than a mere tinkering with details. The theory is fundamentally flawed. I noted a lack of fit between two of the ideas that are near and dear to Lloyd's heart, namely, reliability of object detection and the use of multiple channels, then reviewed the prospects Lloyd has of patching the theory here. In addition, I argued that multiple channels and conditional probabilities do not display the necessary selectivity to handle the linear and lateral selectivity problems. Third, I argued that the dialectical theory of errors of commission is inadequate.

Taking a broader view of correlational semantics, we might recall that all theories have problems, but that some theories have more problems than others. Fodor's asymmetric causal dependency theory has its problems (Adams and Aizawa, 1992, 1993, forthcoming; Cummins, 1989; Maloney, 1991), as does Dretske's early information theoretic account (Dretske, 1981; Fodor, 1990b). Still, Lloyd's theory may have more than its fair share of problems. At the very least, Lloyd would seem to face considerable odds in his next instalment on this theory.

Department of Philosophy
Central Michigan University
Mt. Pleasant, MI 48859
USA

References

Adams, F. and Aizawa, K. 1992: 'X' Means X: Semantics Fodor-Style. *Minds and Machines*, 2, 175–83.

Adams, F. and Aizawa, K. 1993: Fodorian Semantics, Pathologies, and 'Block's Problem'. *Minds and Machines*, 3, 97–104.

Adams, F. and Aizawa, K. forthcoming: Fodorian Semantics. In T. Warfield and S. Stich (eds), *Mental Representation: A Reader*. Oxford: Blackwell.

Adams, F., Aizawa, K. and Fuller, G. In preparation: A Cure for Disjunctivitis.

Churchland, P. 1979: *Scientific Realism and the Plasticity of Mind*. Cambridge University Press.

Cummins, R. 1989: *Meaning and Mental Representation*. Cambridge, MA.: MIT Press.

Dretske, F. 1981: *Knowledge and the Flow of Information*. Cambridge, MA.: MIT Press.

Dretske, F. 1986: Misrepresentation. In R. Bogdan (ed.), *Belief*. Oxford University Press.

Dretske, F. 1988: *Explaining Behavior*. Cambridge, MA.: MIT Press.

Fodor, J. 1987: *Psychosemantics*. Cambridge, MA.: MIT Press.

Fodor, J. 1990a: *A Theory of Content and Other Essays*. Cambridge, MA.: MIT Press.

Fodor, J. 1990b: Semantics Wisconsin Style. Reprinted in Fodor, 1990a.

Lloyd, D. 1987: Mental Representation from the Bottom Up. *Synthese*, 70, 23–78.

Lloyd, D. 1989: *Simple Minds*. Cambridge, MA.: MIT Press.

Maloney, C. 1991: Mental Misrepresentation. *Philosophy of Science*, 57, 445–58.

Salmon, W. 1970: Statistical Explanation. In W. Salmon (ed.), *Statistical Explanation and Statistical Relevance*. Pittsburgh: University of Pittsburgh Press.

Stampe, D. 1979: Toward a Causal Theory of Linguistic Representation. *Midwest Studies in Philosophy*, 2, 81–102.