Anomaly and Quantification

James R. Shaw
Draft, please don't cite

In this paper, I'll argue for several theses about *semantically anomalous* utterances—utterances of grammatically well-formed sentences that resist interpretation like typical utterances of (1)–(3).[1]

(1) * Quadruplicity drinks procrastination.[2]

(2) * This stone is thinking about Vienna.

(3) * He put an event in the hole.

The way I'll argue for these theses is by exploring the relatively ignored topic of how anomaly interacts with natural language quantification—a topic which has a surprising power to shape our understanding of the nature of anomaly and its projection behavior, and to reveal some unique and unappreciated influences anomaly might have on semantic and logical theorizing.

In §1, I'll begin by arguing that anomaly generates a unique form of semantically enforced quantifier domain restriction. I'll then argue that the best explanation for why anomaly interacts with quantifiers in this way is that anomalous utterances are truth-valueless. In §2, I'll explain why these results yield some novel ways for truth-valuelessness to influence the shape of our semantic theories, including by allowing anomalous material to play an unprecedented 'positive' role in the compositional semantics of truth-evaluable utterances. Lastly, in §3, I'll explore how the interaction of anomaly with quantifiers may impact views in philosophical logic. I'll argue that the ways in which anomaly does not project in quantified contexts generate unique forms of classical inference failure, and that the ways in which anomaly does project in quantified contexts provide special motivations for reconceptualizing our logical consequence relations.

1 Anomaly and the Semantics of Quantifiers

I'll shortly appeal to facts about the interaction between anomaly and quantification to argue for the following claim.

[1] These sentences are sometimes called "category mistakes". I prefer to avoid this terminology since it is connected with a *theory* about semantic anomaly which I take to be erroneous but do not have the space to discuss here: that they are the product of mismatches of logical types or 'categories'.

[2] I'll use the marker "*" to mark general oddity and possible anomalous status.

(C) When an uttered sentence exhibits the resistance to interpretation characteristic of anomaly in a given context, that utterance is not truth-valued.

Before I begin, some remarks on (C)'s formulation: First, though (C) treats utterances as potential bearers of truth, I don't mean to be taking a stand as to what the proper bearers of truth are. I'm only presuming that we may at least derivatively think of utterances as truth-bearers. Second, though (C) is a highly controversial claim, there are several noteworthy respects in which (C) is weak. For example, (C) is non-committal as to whether anomalous status precludes an utterance from expressing a proposition in its context. Whether this is so might turn, for example, on the question of whether trivalence is a property propositions could coherently bear. (C) is also non-committal as to whether anomalous character arises in a context-independent way, and so leaves open that a sentence figuring in an anomalous utterance could be used in other circumstances, or in embedded contexts, truth-evaluably.

In §1.1, I'll argue for (C) indirectly by showing that it best explains the source of a peculiar type of quantifier domain restriction. This argument proceeds in abstraction from whether the domain restriction has syntactic, semantic, or pragmatic origins. In §1.2, I'll present some evidence that the domain restriction is best construed semantically.

1.1 Quantifier Domain Restriction

Consider the following, admittedly uneventful, narrative.

> **Trees and Planks**. Bob owns a house with a large yard. In the yard there are six trees and six beautiful hand-carved Scandinavian planks, but nothing else—no bushes, brush, grass or anything of the sort: just dirt. Bob wants to build a fire to keep warm in the winter but is loathe to use those wooden planks (they were a gift from his mother). Consequently Bob uproots the six trees and uses them as firewood.

In response to such a story, speakers are typically willing to classify (4) as true and (5) as false.[3]

(4) Bob uprooted everything in his yard and burned it.

(5) Bob burned everything in his yard.

[3] As informally tested, speakers make these assessments the vast majority of the time. I don't want to presume that assessments are completely uniform, and the argument I shortly give on the basis of the evaluations of (4) and (5) only requires those evaluations to occur *some* of the time. Near the end of §1.2, I'll discuss some factors which might help explain why judgments would fail to be uniform.

This presents us with a puzzle: (4) should entail (5) as is witnessed in what look to be plausible renditions of their logical form below.

(4′) $[\forall x\text{: InYard}(x, \text{Bob})][\text{Uproot}(\text{Bob}, x) \wedge \text{Burn}(\text{Bob}, x)]$

(5′) $[\forall x\text{: InYard}(x, \text{Bob})][\text{Burn}(\text{Bob}, x)]$

In one sense there is an easy explanation for why the apparent inference is blocked—an explanation transparent to English speakers. In (4) the planks of wood aren't being considered part of "everything" whereas in (5) they are. That is, the domain of quantification in (4) is restricted so as to exclude the planks, while in (5) the domain broadens.

I want to start by focusing on a question about the *source* of the witnessed domain restriction: what brings it about that the planks are removed from the quantifier domain in (4)? This question is independent from the question of whether the domain restriction is a semantic or a pragmatic phenomenon. If the restriction is a semantic phenomenon, for example, the question arises as to why the semantics for (4) is unresponsive to the status of the planks. If it is pragmatic, on the other hand, the question still arises as to why (4) doesn't communicate information about planks.

I think we can best answer this question by considering the anomalous (6).

(6) * Bob uprooted the planks.

I think that the best explanation for why the restriction occurs incorporates two claims: (6) is not truth-valued, and the quantifier domain in (4) is responsive to that form of semantic defect. If this is right, then it would show an instance of (C) to hold. Since similar examples to (4) and (5) can be constructed for other instances of anomaly, this in turn supplies a schematic argument lending piecewise support to (C) itself.

Since my argument involves an inference to the best explanation, I'll begin by examining some alternative accounts of the source of the domain restriction that fail. These alternatives construe the quantifier domain in (4) as responsive either to the salience and relevance of the planks for conversational purposes, or to the the falsity or oddity of (6). These options are not meant to be exhaustive, but representative in a way that will come out in the discussion.

The phenomenon of quantifier domain restriction has already been the subject of thorough investigations by linguists and philosophers. These investigations tend to focus on cases like (7).

(7) All the beer is in the fridge.

You will rarely hear an utterance of (7) used to communicate that a given fridge contains all the beer in the world. Rather, some suitably salient instances of the beverage will be up for discussion—for example all the beer someone bought at the store a given day, or all such beer except the two bottles that person dropped on the way home.

Cases like (7) are of interest because their domains of quantification are highly sensitive to features of conversational context, and so form central cases for investigating general theories about how context interacts with language use. For example, whether (7) is usable to communicate the first or the second of my two readings above might depend essentially on whether it is used in a context in which two bottles broke on the trip home, and this fact is apparent to all parties in the conversation.

The domain restriction in (4), however, exhibits at least two features which distinguish it from cases like (7). First, the domain of quantification shifts from (4) to (5), whereas the context need not alter significantly between their assessments. That is, evaluators of those sentences who read them one right after the other nonetheless make the truth-value attributions indicative of a domain shift. The shift even occurs when the sentences are evaluated in reverse order.[4]

(4) also differs from (7) in terms of its responsiveness to considerations of salience and relevance. Consider again an utterance of (7). One way of explaining why its domain of quantification includes, say, all and only beer bought at the store on the day of the utterance even though there is also some beer stashed in the basement, is that the former beer is more salient than the latter for the purposes of the conversation at hand. Were the other beer made salient enough, and relevant to the topic of conversation, it would fall into the domain of quantification as well, as is witnessed in (8).

(8) You remember that beer that we bought at the store? Well, it turns out there was even more in the basement. And guess what: all the beer is in the fridge.

So typically rendering an object suitably salient or relevant ensures it will be included in the domains of subsequently used quantifiers.

The domain restriction in (4) is not responsive to salience and relevance in the same way.

(9) Bob was cold the other day and looking for kindling to keep warm. The type of trees that grow on Bob's property were not really any good for making fires, but the Scandinavian planks in his yard were spectacularly flammable. Bob didn't really value those planks at all. Anyway, at the end of the day he uprooted everything in his yard and burned it.

It is hard to make the planks more salient than in (9). They are not only clearly up for discussion, but are stressed as *directly relevant* to the topic of the last

[4]Of course, I do not mean to be saying that there is no change in context between utterances or readings of (4) and (5). Indeed, between any two utterances there always is some such change, namely the one produced by the passage of time and the fact that it enters into the conversational record that a new utterance has been produced. But an appeal to such changes to explain the domain restriction in (4) would be a hard sell, especially given the stability of speaker assessments in the reverse ordering.

sentence of the monologue: what Bob burned. Nonetheless they are kept out of the quantifier domain of that final sentence—an instance of (4).

In cases like (7), we can claim that the lack of salience of certain objects, or the fact that those objects are not obviously 'up for discussion' in the conversational context *explains* why those objects fall outside the relevant quantifier domains. This has a good deal speaking in its favor, since it not only conforms to the data, but makes sense of why contextual domain restriction arises: it is part of a strategy to gain efficiency in communication by letting conversational context dictate the bounds of quantifier domains instead of having speakers explicitly delimit them. But since the domain restriction in (4) is not responsive to salience or relevance in the right ways, this explanation for (7) cannot be transposed to (4).

If we can't explain the restriction in (4) in the same way as for (7), what are our other options? A second natural attempt involves claiming that the domain restriction occurs because, were it not performed, the quantified statement would be false as it would have a false instance correlated with (6). As part of some process of avoiding a reading which is so problematic, another similar reading is sought or generated which is not (or not so obviously) false. Perhaps there is a semantic mechanism which restricts the quantifier domain to yield this effect, or something like an application of a principle of charity leads interlocutors to search for the relevant reading.

While the sort of mechanism this option posits is arguably widespread in other contexts, it is just as unpromising an explanation of the quantifier restriction in (4) as the appeal to salience or relevance. This time, although the explanation has the potential to capture the datum given by (4), the general principle it invokes simply does not apply in the grand majority of cases involving quantification over false, or even trivially false, instances.

Suppose, for example, that the story of *Trees and Planks* were modified so that Bob didn't burn one of the trees in the yard. Then speakers immediately, and unproblematically, take (4) to be false. Speakers make no attempt to rectify the utterance by removing the relevant trees from the domain of quantification. And indeed, such a response would be quite bizarre.

Consider another, more forceful example: if one mathematician utters (10) to another mathematician after handing him a page on which the numbers 2 through 5 are written, it is implausible that any domain restriction would be accommodated.

(10) Every number written on that page is prime.

The interlocutor would almost certainly conclude there had been an oversight on the part of the speaker. This is so even if, as we have supposed, the relevant parties are experts and highly unlikely to make the relevant mistakes. It is also the case even if the falsehood in question is both an obvious and a *necessary*, rather than a contingent, falsehood.

There are, of course, cases where something's being obviously false might create a domain restriction *via* the normal contextual forms of quantifier domain restriction. Consider the following case: Clyde runs into a room where Al is standing, grabs a small pile of books on the table, leaving only a pen on it, and runs out. Bill enters the room, sees only the pen on the table, and asks Al what happened. Al might successfully communicate the facts by uttering (11).

(11) Clyde grabbed everything on the table and ran with it.

He might do this, despite the pen being on the table in plain view, counting on Bill to recognize from the context that the pen is not among the things talked about. However, the fact that the restriction proceeds via the normal contextual avenues means that considerations of salience and relevance may defeat it. Suppose Bill enters the room, sees the pen, and utters (12).

(12) I can't believe I left my precious antique pen on the table where Clyde could just take it. I see that Clyde grabbed everything on the table and ran with it. I'm so happy he didn't grab that pen on the table.

Though it is possible to figure out what Bill means, his utterance sounds contradictory. The tension between the fact that Bill's utterance has rendered the pen pertinent to the taking, and the falsity of Bill's quantified statement on its most general reading is very readily felt.

Thus the falsity of certain substitution instances—even their obvious, egregious, and necessary falsity in the face of mutually aware interlocutors—does not in general produce a domain restriction. False instances may restrict quantifier domains indirectly via normal modes of contextual quantifier domain restriction, but in this case the domain restriction will ultimately be very sensitive to considerations of salience and relevance. And, as we have already seen, any type of domain restriction which is sensitive in these ways is not the type we are out to explain.

Both of the first two attempted answers to the question concerning the domain restriction occurring in (4) suffer from an obvious defect: they fail to take into account that the domain restriction looks to be connected with the *anomalous* character of the substitution instance, (6), which engenders the restriction in (4). That anomalous character is almost certainly the product of *some* kind of defect. The first account I examined, which transposes the account of the domain restriction in (7), does not do any justice to the idea that the domain restriction in (4) is responsive to the semantic properties of its constituents, rather than simply to features of the conversational setting. The second account, which took the alleged falsity—perhaps the egregious, obvious, and necessary falsity—of (6) as the grounds for the restriction, at least looked to some semantic properties of the sentence to track the relevant restriction. The problem is that it didn't look to anything that was specially connected with anomaly. To that extent, it ended up positing a general principle of domain restriction which was simply not witnessed in cases where anomaly was absent.

A third possible explanation takes its cue from these failures, and acknowledges that there is something particular about anomaly which helps enforce the domain restriction. On this account, there are just things which it is 'strange' or 'awkward' to talk about using certain predicates. Planks, for example, exhibit this awkwardness as concerns the predicate "uproot". Our judgments of anomalous status tend to track this strangeness. Moreover, it is this strangeness of predications, and not necessarily their falsity, which results in a domain restriction. Either a semantic mechanism is in place to track these peculiarities, or speakers interpret their interlocutors charitably by not imputing awkward claims to them.

This third explanation faces problems, not in that it is inadequate, but in that it is underspecified. What exactly makes a particular ascription of a predicate awkward or strange in the relevant sense? The more specific one is about what the peculiarity consists in, the more implausible a domain restriction over 'awkward' substitution instances becomes. Here, for example, are three possible ways of spelling out the requisite kind of strangeness.

(a) Predicating F of a is strange if people tend not to make such predications, or tend not to be moved to make them.

(b) Predicating F of a is strange if it describes a highly fantastical or wondrous situation.

(c) Predicating F of a is strange if it is particularly confusing or difficult to understand.

I won't dwell on any of these elaborations, since it should be fairly clear that there is no such general domain restriction over the members of any of the classes described by (a)–(c). Indeed, one would hope that there would be no such domain restriction, since this would seem to lead to a strange and completely unnecessary form of expressive limitation in natural language.

Of course, there are other ways of spelling out the 'strangeness' alluded to in the third response. The point is that simply saying that it is strangeness which generates the domain restriction passes the buck in a theoretically unsatisfying way, and that most ways of spelling out what the strangeness is will be *too general*. Any elaboration threatens, as with (a)–(c), to predict far more instances of domain restriction than there actually are.

Our discussion so far narrows our options quite a bit and helps point to an obvious solution. We need an explanation for why there is a salience-insensitive form of domain restriction which occurs over anomalous substitution instances—so it should draw on something that is special about those instances. Appealing to the claim that anomalous sentences are not, in general, truth-evaluable furnishes us with the materials to provide such an explanation by supplying a distinct semantic feature specially borne by anomaly.[5]

[5] I'm not claiming that anomalous utterances are the only ones to fail to be truth-evaluable. In saying that failures of truth-evaluability would 'specially' be borne by anomaly, I just mean if anomaly was truth-valueless, this would distinguish it from obvious falsehoods, 'odd', 'fantastical', or 'confusing' claims, and other kinds of claims which fail to generate the kind of quantifier domain restriction I've been focusing on.

Taking anomalous sentences to be truth-valueless is only half of the story, though. We also need to explain how this truth-valueless status interacts with quantifier domains. The basic idea is that *the domain restriction is the product of a general interpretive strategy on the part of language users to maximize truth-evaluable content*. This strategy has implications for the semantics of quantified statements with anomalous instances on the following plausible claim: that some uttered quantified statements *would* be truth-valueless were their domains of quantification to include *all* of their truth-valueless substitution instances. Some evidence for this claim, on the assumption of (C), comes from basic projection data for anomaly. For example, sentences with quantifiers uninterpretable unless forced to range over anomalous instances, such as "the number eight is red" or "all towels are prime", tend to be anomalous. The claim also gains motivations from many standard trivalent semantics. For example, the claim will be made true by some uses of universal quantification on both the Strong and Weak Kleene schemes.

If some quantified statements inherit truth-valuelessness from some of their truth-valueless substitution instances, and (C) is true, then a policy of restricting their domains of quantification to exclude relevant anomalous substitution instances preserves the truth-evaluability of many whole quantified statements, leading to a straightforward increase in conventional expressive power.

To see what I mean by this, suppose (C) holds and consider maximally general assertions of various kinds. Take, for example, the wise words of the Duchess to Alice.

(13) Everything has a moral if only you look for it.

Now what was the Duchess talking about? Presumably books, fables, tall-tales, but also (why not?) the lives of great men and women and incidental events in one's own daily life. Note, however, that there are of course things which the Duchess clearly is not talking about: tea doilies, bowling alleys, and socks among them. To say of these things that they have a moral would be anomalous. If these instances are truth-valueless and would render a fully general interpretation of (13) truth-valueless as well, then we have an explanation for why it might be advantageous to restrict the domain of quantification in (13) over non-anomalous objectual substitution instances. This would enable (13) to serve as a *truth-evaluable* (and quite possibly true) statement.

If, by contrast, we allow quantification to range over all objectual substitution instances, we would be in danger of condemning many generalities like (13) to fail to express truth-evaluable propositions relative to a context—a strict loss of conventional expressive power with no corresponding gain. Moreover, as the Duchess' example hopefully makes clear, there is value to being able to express the content of the corresponding generalities with restricted domains.[6]

[6] I speak of 'conventional' expressive power since the truth of (C) and its application here might still allow for the expression of 'unconventional' trivalent propositions. I am presuming that even if there were such entities, they would be in some ways potentially detrimental to communication.

If the hypothesized policy of restricting some quantifier domains to exclude certain truth-valueless substitution instances were in place, it would mostly only be apparent in cases like (13) which can, on their face, be accounted for equally well by standard explanations of contextual quantifier domain restriction. Cases like (4), where salience plays no role, signal that a different phenomenon is at work.

In this way, (C) affords us a plausible explanation of why we witness the special kind of domain restriction that anomaly produces—something which none of the other three responses were even able to achieve. But this is not all. Adopting (C) has three additional virtues. First, as we have seen, it gives an intuitive explanation of why the domain restriction occurs which shows it to be a *communicatively beneficial* linguistic mechanism. Second, the explanation validates what seems to be an obvious fact about the domain restriction: it is a kind of response to a semantic feature which anomaly in particular bears. Finally, it also helps explain why the domain restriction is robust in the face of considerations of salience and relevance. Regardless of how salient a given object is for the purposes of the discussion at hand, this won't change the fact that the respective anomalous instance fails to express a truth-evaluable proposition. A speaker who focuses attention on an object which figures as an anomalous substitution instance of a subsequent quantified statement cannot simply be making a blatant mistake about what the facts are, the way the mathematician who uttered (10) could be—there *is* no mistake about how the world is to be made. Thus there is no immediate cause to reinterpret their statement as having a quantifier which ranges more broadly.

If this is right, we have strong support for (C). It figures as an indispensable part of the best account of the source of the unique form of quantifier domain restriction witnessed in (4) and similar cases of domain restriction from anomaly. For this reason, I propose that we accept (C) and explore what changes the domain restriction which motivated it might require of our compositional semantic theories. Before tackling this issue in §2, however, I need to address the question as to whether or not the domain restriction from anomaly is semantically enforced, and how prevalent it is.

1.2 Syntax, Semantics, Pragmatics

The argument of the previous subsection for the claim that anomalous utterances are not truth-valued proceeded in abstraction from the question of whether or not the domain restriction from anomaly was enforced syntactically, semantically, or pragmatically. One could, for example, accept (C) on the basis of the arguments I gave and still maintain any of these three options: the domain restriction could be the product of phonetically null syntactic material, it could be the result of a systematic semantic mechanism, or perhaps a statement like (4) should be counted as straightforwardly truth-valueless with pragmatic principles leading speakers to reinterpret the utterance in the appropriate way. In this section, I'll argue that the restriction is best construed as operating at the semantic level. This claim will play an important role in §2 and §3.

First, just as in cases of salience-sensitive quantifier domain restriction, maintaining that the domain restriction in (4) is syntactic is difficult due to the problem of underdetermination.[7] Very few, if any, theorists think that the domain restriction in (7) is the result of added, unarticulated syntactic material.

(7) All the beer is in the fridge.

The reason is that there seems to be no principled way to pick out one of many extensionally equivalent expressions allegedly present in a given utterance of (7) to restrict the quantifier domain appropriately. For example, in one particular context "all the beer *which we just bought*" does just as well as "all the beer *which we just bought today*", and "all the beer *which we carried in together*" and so forth. In the same way there are extensionally equivalent ways of bringing out the domain restriction in (4): "which has roots", "which is rooted in the ground", "which is planted in the yard", and so on. In addition to the problem of selecting one from among *several* possible candidates for syntactic ellipsis, sometimes it is difficult to even find one. Consider again, in this regard, the Duchess' (13).

(13) Everything has a moral if only you look for it.

On its maximally general interpretation, which only excludes anomalous instances, (13) is restricted to a highly diverse array of things. Even if one could find the right syntactic material to perform the restriction, it would have to be implausibly long, and completely unavailable to the speakers supposedly generating the relevant syntactic structure.[8]

The real question is whether the phenomenon is semantic or pragmatic in nature. Three things point to a semantic treatment. First, the domain restriction, where it occurs, is fairly robust: it is salience-insensitive and gives evidence of being quite systematic. Second, unlike with more familiar forms of contextual quantifier domain restriction, the domain restriction due to anomaly is triggered by the presence of a *semantic feature*—anomalous status—rather than by special features of the context of use. As I've allowed above, whether or not an utterance *is* anomalous may turn out to be sensitive to conversational context. But within a fixed context, judgments of anomalous status are clearly tracking some kind of semantic aberration. It seems reasonable to suppose that interpretive shifts clearly responsive to the presence of a semantic feature are themselves semantic. Both of these first two reasons for treating anomaly semantically are connected with the fact, which I'll explore in §2, that we can systematize the information relevant to the domain restriction so as to track when it occurs.

[7]The basic argument can be found in Stanley & Szabó (2000a) p.236*ff.*.

[8]These considerations against what I am calling the syntactic mode of quantifier domain restriction naturally do not speak directly against the idea that there could be, say, a demonstrative or a variable in the syntactic structures of these quantified statements which pick up their semantic values from other elements in the sentence. I'll discuss this possibility in §3.

A third reason for treating the domain restriction from anomaly semantically is that there are noteworthy arguments for considering even the more seemingly pragmatic phenomenon of domain restriction from salience or relevance as semantic in nature. These can be extended to apply to the domain restriction from anomaly. The arguments, owing to Stanley & Szabó (2000a), focus on a kind of *binding* phenomenon arising from the interaction of multiple quantifiers. For example, on its most natural interpretation, the domain of quantification of "every bottle" in (14) varies with the different rooms in the domain of the first quantifier.

(14) In every room in John's house, every bottle is in the corner.

Stanley and Szabó have argued that pragmatic accounts of how quantifier domains are restricted can have a hard time explaining how the natural reading is arrived at in (14) since there appears to be a kind of binding.[9] (14) reads roughly as "every x such that x is a room in John's house is such that every bottle *in* x is in the corner *of* x." Semantic accounts which posit a variable that interacts with the discourse context to produce contextual domain restrictions can account for this phenomenon very easily, since the binding in question can occur over the relevant, syntactically realized variable.

I don't want to take a stand on whether Stanley and Szabó are right. I only want to note that *if* they are right, then this provides additional support for treating anomalous domain restriction semantically. This is because the domain restriction from anomaly can *interact* with this binding phenomenon.

> **Jose's Cocktails**. At a gastronomical competition Jose served three courses, the latter two accompanied by different rum cocktails. After sipping from the cocktails, the judges declared that both would benefit from the addition of tiny amounts of select spices. The judges accordingly added four different spices, two to each drink, and gave them back to Jose to taste.

(15) The judges sipped from everything Jose served before adding two spices.

Speakers take (15) to be true of the above story (as if the sentence read "... every *drink* Jose served..."). For (15) to be true, the pairs of spices talked of in (15) must be relativized to the objects Jose served. But the elements quantified over by "everything" are said to be "sipped" and so, according to the present view, should be restricted to the objects capable of standing in that role. This forces speakers to exclude the meals from the domain of that quantifier, contributing to the true reading. But *if* there is binding of "two spices" by the quantifier "everything" at the semantic level, then to get the true reading we need the values of the relevant bound variable to be restricted as well in order to have

[9] They have some recourses though, at least by appealing to particular theories of structured propositions, as Stanley & Szabó (2000b) concede.

the appropriate pairs of spices talked about. This would occur much more naturally if the restriction due to anomaly were processed at the semantic level as well. Otherwise, for example, at the semantic level "two spices" is bound by a variable from an unrestricted quantifier, requiring the explanation of the appropriate bound reading to be significantly more complex. To get the true reading, the binding and the domain restriction from anomaly must be operating in tandem.

What I've been arguing so far is that the case for taking anomalous domain restrictions to be semantic in nature is strong—significantly stronger, for example, than the case for taking contextual quantifier domain restriction to proceed via a semantic influence of context. We've just seen that the best reasons in favor of the latter case, including the binding phenomenon, also favor construing anomalous restriction semantically. Moreover, unlike with contextual quantifier domain restriction, the insensitivity of anomalous domain restriction to shifts in contextual salience, and the fact that it arises from a responsiveness to a semantic feature both further motivate giving it a systematic semantic treatment. For these reasons, I'll proceed now on the working assumption that the domain restriction from anomaly is semantically enforced to explore what consequences this has for semantic theorizing and logic.

But before I continue, I need to mention a brief caveat. When I claim that the domain restriction is semantic, I do not mean to claim the phenomenon exhibits no sensitivity to context, or is exceptionless. It is not necessarily insensitive to context because anomalous status itself may be context-sensitive. It need not be exceptionless because, if the story of §1.1 is on track, the restriction is merely a default interpretive mechanism, which may be overcome by other factors. This is important because there is evidence the domain restriction is not entirely uniform. I suspect, however, the explanations for the lack of uniformity are diverse. For example, processing costs may be relevant. Speakers might be less likely to enforce the domain restriction if the anomalous triggering material occurs late in a sentence, past the point where the ordinary material contributing to the quantifier restrictor occurs as in (16).

(16) Bob burned everything in his yard with due precaution, not long after having uprooted it.

Also, there are degrees of anomalousness, and there is an open question of 'how anomalous' something must be to count as truth-valueless, and generate the domain restriction. Finally, there is evidence that anomalous status itself might be a context sensitive matter, as I've already stressed. Unfortunately I don't have the space to explore the interactions between these phenomena and the domain restriction here. What's important is that the view I mean to be defending leaves open that that the domain restriction could have exceptions for these reasons and perhaps others.

2 Semantics for Anomaly

If, as I have argued, the quantifier domain restriction of §1 is a semantically enforced phenomenon responsive to the truth-valuelessness of anomaly, what can we learn from this about the shape of our compositional semantic theories? First, in §2.1, I'll make some general remarks about what I take to be a plausible answer to this question. Then in §2.2, I'll sketch a formalism which encapsulates the ideas that arise in the discussion.

2.1 Projection and Contributions to Interpretability

Anomalous sentences are semantic aberrations and their defects ought to be captured by a semantic theory. Here, for example, is a strategy that will not work: treating anomaly as malformed in the same way as ungrammatical sentences, thereby excluding them from the purview of one's formal semantic apparatus. That would be a bad strategy, since the semantics will ultimately need to retrieve information about which sentences are anomalous in order to adequately track the quantifier domain restriction due to anomaly. Syntax must admit the anomalous sentences if the semantic apparatus is to systematize the relevant information.

Once anomalous sentences are treated by the semantic theory, the latter will naturally have to make special provisions on pain of misrepresenting the semantic values of anomalous and non-anomalous sentences alike. For example, purely bivalent model theoretic semantics—by which I mean semantics which only have two possible extension assignments for whole sentences—have no resources to mark off the distinction between anomalous sentences and those which are simply false, as the theorems entailed by the theory make a bipartite division of sentences relative to a model-index-context triple $\langle \mathcal{M}, i, c \rangle$ (say).[10]

The way around the problem is simple: add a semantic status over and above truth and falsity and use this value to recursively track instances of anomaly and their effect on quantifier domains. The idea of accommodating a third semantic status is not a particularly new one—even as applied to anomaly.[11] Nonetheless the move to trivalence raises several important foundational questions for semantics that sometimes go unappreciated. One such question concerns the added information taken for granted at the level of lexical semantics—information, for example, about what objects a given predicate truth-evaluably applies to. Adverting to this information raises interesting questions about whether lexical semantics models only information or competences speakers have, or also information and competences they *lack*. Those questions can only be fully answered by supplying accounts of the nature of the relevant third status—here truth-valuelessness. I won't be pursuing any such account here.

[10] I mean to include "undefined" as a possible extension.

[11] See, for example, Thomason (1972) and Lappin (1981) for trivalent treatments of anomaly.

Instead, I want to focus on a second underappreciated possible consequence of the shift to trivalence: the third semantic value *may play a systematic, positive role in truth-evaluable interpretation.* I have argued that precisely this happens in §1, by arguing that anomalous, and hence truth-valueless, status systematically influences the domains of quantifiers. This idea has the power to shape answers to other foundational questions about the shift to trivalence like that just mentioned. But before I can say more clearly what I mean by a 'positive' role in truth-evaluable interpretation, and why this positive role is unique to the phenomena explored in §1, I need to say a little more about the compositional behavior of anomaly and the domain restriction it creates.

Let me then quickly sketch some points about the tools we need to recursively track the influence of anomaly in interpretation, some of which should be familiar. At the base level of the recursion the semantics will clearly need, for each predicate in the language, information about which objects that predicate can truth-evaluably be used to talk about. A natural step to take here—one which can be found instantiated in many broadly trivalent systems—is to take this information for granted by associating with each predicate a set of objects comprising what I will call, as a nod to Russell, its *domain of significance.*[12] (Note: this is essentially just a terminological variant of the equivalent strategy of adding information about anti-extensions of predicates to information about their extensions.) Since I won't have time to address issues concerning the structure of those domains here, I will presently take the most *non-committal* formalization of them possible: one allowing for *arbitrary* sets of objects, or n-tuples of objects for predicates of higher adicity than one.[13]

Such domains constitute the information needed to classify which 'atomic' ascriptions of a predicate to an n-tuple of objects are anomalous. Using these domains, perhaps along with the extensions of predicates, to recursively track how anomaly projects into coordinative constructions is simple to achieve, though controversial in precisely how to implement. Anomaly seems to have an 'infectious' character: wholes with anomalous parts tend to be anomalous as seen in (17)–(19).

(17) * Relapses demote the undertow.

(18) * Relapses demote the undertow and ice cream tastes great.

[12] See, e.g., Russell (1908), Carnap (1937), and Fodor & Katz (1964) for examples or, again, Thomason (1972) and Lappin (1981) for applications to anomaly. A note on terminology: I prefer to avoid the term "category" for the reasons alluded to in n.1. Also, one shouldn't read too much into the nomenclature. Being outside a predicate's domain of significance should not be interpreted as something to which it is not 'meaningfully' applied in any sense that would come into conflict with the weakness of my earlier claim (C).

[13] This contrasts, for example, with first delimiting sets of logical sorts or categories, and associating an n-ary predicate with an n-tuple of such sorts. Such a theory is more restrictive and I worry, for that reason, may lead to incorrect predictions.

(19) * If relapses demote the undertow, I'm going on vacation.

However, anomaly may not always project in these ways. Possible exceptions include embeddings of anomaly under some uses of negation, and into disjunctions and counterfactual conditionals. My goal here is not to take a stand on how anomaly projects in any of these cases. It suffices to say that many different ways of tracking truth-valueless projection behavior across simple connectives are already well-explored (for example, in Kleene's Strong and Weak schemes). Moreover, these can all be implemented in a framework drawing only on domains of significance (and extensions) for atomic predicates, and so don't bear directly on the novel techniques I want to introduce to cope with the case that has preoccupied me here, namely that of quantification.

So let me turn to the question: when is a quantified statement anomalous? The answer to this question will turn out to be helpful in deciding how to represent quantifier domain restriction in non-anomalous cases. Consider the following two simple instances of quantified anomaly.

(20) * Some primes are red.

(21) * Every tomato is polarized.

From both a logical and a semantic perspective (21), for example, doesn't look so bad. A standard rendition of the logical form of (21) might be as "[$\forall x$: Tomato(x)][Polarized(x)]". Certainly it should be permissible to predicate a variable with "Tomato". Likewise for "Polarized". So, if anything, something must go wrong at the level of appending the quantifier. Similarly (20) merely asserts a non-empty intersection between two sets—the set of numbers and the set of red things. But the intersection of those sets *is* empty. Why isn't the claim simply false?

On reflection there seems to be a ready explanation for why both (20) and (21) are anomalous, brought out by consideration of similar instances of anomaly. There are plenty of primes and plenty of red things, but nothing non-anomalously talked of as both. That is, for any object o either "o is prime" or "o is red" is anomalous. Likewise for tomatoes and polarization. What renders quantified anomalous sentences problematic appears to be non-intersective domains of significance. Otherwise put, a quantified sentence is anomalous if it has only anomalous objectual substitution instances.

If we were only concerned about these kinds of instances of quantified anomaly, how should we track their occurrence? The answer is simple: use domains of significance to recursively keep tabs on which assignments to unbound variables result in a truth-evaluable whole. We can call this set of variable assignments the *domain of significance of the open formula.* If the domain of significance of an open formula is empty, this means any way of binding its variables results in an anomalous substitution instance.[14]

[14]On the semantics I'm sketching the *only* anomalous quantified utterances are those appended to open formulas which taken together have an empty domain of

To complete our semantics for quantified statements we also must at least pronounce on the interpretation of non-anomalous uses of quantifiers as well. This is a more delicate issue since, according to the view I defended in §1, we interpret quantified statements by restricting their domains to avoid having to interpret them in ways that *would* prevent them from being truth-evaluable. For some quantifiers, there are relatively uncontroversial ways of construing how failures of truth-evaluability sometimes project given unrestricted quantification. For example, most trivalent semantics for universal quantification (including both Strong and Weak Kleene, for example) have that quantifier inherit truth-valuelessness from *any* truth-valueless substitution instances within its quantifier domain, provided the other instances are true. Quantifiers like "some", and "most" present us with more controversial options. Will "some Fs G" always exhibit truth-valuelessness just so long as at least one substitution instance from an object in its quantifier domain is truth-valueless? Some semantics—like a Strong Kleene semantics—say "no", while others—like a Weak Kleene semantics—say "yes". Similar remarks apply to "most".

Different proposals for the projection behavior of truth-valueless instances in unrestricted quantifier domains will interact with my proposed views on quantifier domain restriction due to anomaly to generate different empirical predictions. This leads to two virtues. First, my proposal is flexible: it can accommodate any view about truth-valueless projection not only for connectives, but for unrestricted quantifiers. Second, my proposal opens up the possibility of *working backwards* from facts about quantifier domain restriction to claims about anomalous status and projection in unrestricted quantifiers. This methodology is quite useful, since intuitions about domain restriction are often much more stable than intuitions about anomalous status, or its projection.

To see this second virtue in operation, here's an application of the 'working backwards' methodology to "most".

> **Vera's Patient**. Vera has one of her patients, Marla, begin their therapy session by producing drawings and text on a single sheet of paper. Marla scrawls a dozen or so images and writes out the first ten words that come to her mind. Vera picks up the paper and, after reading the first two words in her head, reads the next eight, which seem more significant, out loud to Marla.

Consider:

(22) Vera read most things Marla scrawled on the page out loud.

Speakers tend to read (22) as true, even when the images on the page are stressed as scrawled on it. This would only be predicted, given the views I've

significance. I'm open to weakening this requirement and allowing other forms of quantified anomaly. This concession is connected with my proposal is that there is a default interpretive strategy speakers employ in restricting quantifier domains to preclude anomalous, and hence truth-valueless, status, which might well be 'overridden' by other factors.

been articulating, on the assumption of two facts: it is anomalous to say of Marla's drawings that they are read, and the semantics for "most As B" should treat it as truth-valueless when there are objects o in its quantifier domain such that Bo is truth-valueless.

So to reiterate, any view about projection of truth-valuelessness in unrestricted quantifier domains can be integrated with my proposed views on domain restriction, so there is no need to make general commitments as to what the original projection behavior is, and it is my preference here not to do so. Also, for this very reason, the theory can actually be used to test views about projection from quantifier to quantifier—perhaps with results at times more helpful than tests that appeal to intuitions about truth-valueless claims.

Now, since a recursive characterization of domains of significance for open formulas of the kind I mentioned earlier would implicitly contain information about the truth-evaluability of various substitution instances, that same recursive characterization has all the information needed in our definitions of the truth-conditions of quantifiers to adequately capture the quantifier domain restriction due to anomaly. Truth or falsity of a quantified statement is ascertained by evaluating objectual substitution instances of the open formulas over which it quantifies that would not lead to failure of truth-evaluability were the quantifier domain to include them. So however we negotiate the details of the quantifier domain restriction, recursively characterized domains of significance for open formulas will be necessary not only to track the presence of anomaly in quantified statements, but to interpret non-anomalous, truth-evaluable quantified statements. They will also be sufficient for both purposes.

This concludes my sketch of a semantics for anomaly: add domains of significance for predicates, use this information to recursively track the presence of quantified anomaly, and also to characterize the interpretation of non-anomalous quantification. The most important, and novel, part of my proposal comes of course in the third step. For the more formally inclined, a more detailed implementation of these ideas can be found in the next subsection.

For now, though, I want to return to my earlier claim, which I also billed as having some importance for a foundational understanding of trivalent semantics: that truth-valuelessness can play a positive role in truth-evaluable interpretation. To understand what I mean by this, and why it arises specially from the existence of the quantifier domain restriction due to anomaly, consider what happens in various trivalent semantics under an *expansion* of a predicate's domain of significance—the addition of some objects to a predicate's domain of significance which antecedently lay outside it. In standard trivalent semantics, producing such an expansion *never produces a shift between truth-evaluables.* That is, such an expansion never changes a true claim to a false claim, or a false claim to a true one.

For example, consider the Weak Kleene scheme, on which truth-valuelessness always projects through connectives and quantifiers. If any sentence S changes its truth-value after the expansion of a domain of significance of a predicate P with an object o, it must be because the truth-value arising from predicating P of o matters to the truth-value of S. In the Weak Kleene scheme the only way predicating P of o could have originally influenced the truth-value of a sentence S is *by rendering it truth-valueless*. So no expansion of a domain of significance on this semantics moves us from a true claim to a false claim, or a false claim to a true claim. The same is true of both the Strong Kleene scheme and a supervaluation scheme.

This commonality isn't incidental. On a prevalent, usually tacit, assumption reflected in trivalent semantics of many kinds—especially those applied to anomaly—truth-valuelessness *merely interferes* with truth-evaluable interpretation. Only the extent to which it actually does interfere is contested. On the Weak Kleene scheme it interferes as much as possible. On the Strong Kleene, less so. And on a supervaluation scheme, less still.

This changes once we move to a semantics which integrates an involvement of truth-valuelessness in producing a quantifier domain restriction, for the obvious reason. Sometimes a statement in which the domain restriction takes place, like (4), has the conventional truth-value it does (true, in this case) because a particular predication is truth-valueless. If we could expand the domain of significance of "uproot" to include planks, we might well change the value of (4) to false.

What this means it that on this semantics truth-valuelessness isn't merely interfering with conventional, truth-evaluable interpretation. It's *contributing* to it. Language users attend to anomaly as they try to figure out what they are conventionally, and successfully saying to each other. As I noted earlier this has the potential to influence our understanding of foundational questions about the character of a third semantic status. This is most obviously so for my earlier question about whether a lexical semantics incorporating a third value in application to anomaly is modeling competences speakers have or competences speakers lack. Other questions may be influenced as well, but unfortunately I won't be able to pursue them any further here. For now, let me give the sketched semantics here a more rigorous formal implementation. Then I will turn to some issues in philosophical logic which are also affected by the interaction between anomaly and quantification.

2.2 Formalism

What follows is an idealized trivalent model-theoretic semantics for anomaly of the kind just described, assuming the simplest projection scheme for connectives and quantifiers: the Weak Kleene scheme. This subsection can safely be skipped if desired.

I'll work in a language incorporating two binary quantifiers $\forall$ and $\exists$, and connectives $\neg$, $\wedge$, $\vee$ with the following syntax.[15]

[15]Here and in the foregoing I've omitted discussion of functions for the sake of brevity. I'm also a bit loose on use mention distinctions.

Syntax

Given

a set of *constants* $\mathcal{C}$,

a set of *variables* $\mathcal{V}$, and

a set of *relation symbols* $\mathcal{R}$ of various adicities:

All $c \in \mathcal{C}$, $v \in \mathcal{V}$ are *terms*.

All $R(\tau_1, \ldots, \tau_n)$ for n-adic $R \in \mathcal{R}$ and terms τ_i are (*atomic*) *formulas*.

If ϕ, ψ are *formulas*, $v \in V$, then $\neg\phi$, $\phi \wedge \psi$, $\phi \vee \psi$ and $(\forall v : \phi)(\psi)$, and $(\exists v : \phi)(\psi)$ are *formulas*.

To the standard characterization of a model, we need to add only information about which predicates truth-evaluably apply to which objects. I use a double-bracket notation $\langle\!\langle\ \rangle\!\rangle$ for such domains. Otherwise the definition of a model is the familiar one.

Models

A *model* $\mathcal{M}$ is a tuple $\langle M, \mathcal{I}\rangle$ consisting of a non-empty *universe of discourse* M, and an interpretation function $\mathcal{I}$ which maps:

– each constant $c \in \mathcal{C}$ to an element $[\![c]\!]^{\mathcal{M}}$ of M;

– each n-ary $R \in \mathcal{R}$ to a pair $\langle [\![R]\!]^{\mathcal{M}}, \langle\!\langle R\rangle\!\rangle^{\mathcal{M}}\rangle$ containing

(i) an *extension* $[\![R]\!]^{\mathcal{M}} \subseteq M^n$; and

(ii) a *domain of significance* $\langle\!\langle R\rangle\!\rangle^{\mathcal{M}}$ with $[\![R]\!]^{\mathcal{M}} \subseteq \langle\!\langle R\rangle\!\rangle^{\mathcal{M}} \subseteq M^n$.

Denotations are also computed in standard fashion. Constants acquire their denotation from the model and variables from a variable assignment.

Denotation

Let $\mathcal{G}$ be the set of *assignments*, i.e. of total functions from $\mathcal{V}$ into M.

The denotation of a term τ in a model $\mathcal{M}$ under an assignment $g \in \mathcal{G}$, written $[\![\tau]\!]^{\mathcal{M},g}$ is given by:

$[\![v]\!]^{\mathcal{M},g} = g(v)$ for $v \in \mathcal{V}$

$[\![c]\!]^{\mathcal{M},g} = [\![c]\!]^{\mathcal{M}}$ for $c \in \mathcal{C}$

With denotations in hand we can specify the domains of significance for more complex expressions. Recall: the domain of significance of an open formula is the set assignments of objects to free-variables in that expression which would result in truth-evaluability.

In making this definition, as I said, I'll use a Weak Kleene characterization of the behavior of $\neg$, $\wedge$, and $\vee$. As noted in the previous section, nothing important hangs on this choice for present purposes (aside from simplicity of exposition), and we can let empirical considerations guide our selection of projection schemes. Quantified formulas, however, merit special commentary. When a formula is appended with a quantifier binding a variable v we need two effects on the domain of significance. First, if every assignment to the unbound variables in the formulas comprising the quantifier restrictor and matrix makes at least one non-truth-evaluable, the quantified formula should inherit this defect. Otherwise the quantified formula itself will be truth-evaluable according to the criteria I've given, and so should itself have a non-empty domain of significance. Now that v is bound, though, the assignments in the formula's domain of significance should be 'indifferent' to the value on v. Both effects are achieved by appropriately importing the domain of significance of the restrictor and matrix and allowing its assignments to be arbitrarily permuted on v.[16]

Domains of Significance

The *domain of significance of an expression e in a model $\mathcal{M}$*, written $\langle\!\langle e \rangle\!\rangle^{\mathcal{M}}$, is a subset of $\mathcal{G}$ given as follows:

$\langle\!\langle v \rangle\!\rangle^{\mathcal{M}} = \mathcal{G}$ for $v \in \mathcal{V}$.
$\langle\!\langle c \rangle\!\rangle^{\mathcal{M}} = \mathcal{G}$ for $c \in \mathcal{C}$.
$\langle\!\langle R(\tau_1, \ldots, \tau_n) \rangle\!\rangle^{\mathcal{M}} = \{g \in \mathcal{G} \mid \langle [\![\tau_1]\!]^{\mathcal{M},g}, \ldots, [\![\tau_n]\!]^{\mathcal{M},g} \rangle \in \langle\!\langle R \rangle\!\rangle^{\mathcal{M}}\}$
for atomic formulas $R(\tau_1, \ldots, \tau_n)$.
$\langle\!\langle \neg\phi \rangle\!\rangle^{\mathcal{M}} = \langle\!\langle \phi \rangle\!\rangle^{\mathcal{M}}$
$\langle\!\langle \phi \wedge \psi \rangle\!\rangle^{\mathcal{M}} = \langle\!\langle \phi \rangle\!\rangle^{\mathcal{M}} \cap \langle\!\langle \psi \rangle\!\rangle^{\mathcal{M}}$
$\langle\!\langle \phi \vee \psi \rangle\!\rangle^{\mathcal{M}} = \langle\!\langle \phi \rangle\!\rangle^{\mathcal{M}} \cap \langle\!\langle \psi \rangle\!\rangle^{\mathcal{M}}$
$\langle\!\langle (\forall v : \phi)(\psi) \rangle\!\rangle^{\mathcal{M}} = \{g[v \to m] \mid g \in \langle\!\langle \phi \rangle\!\rangle^{\mathcal{M}} \cap \langle\!\langle \psi \rangle\!\rangle^{\mathcal{M}}, m \in M\}$
$\langle\!\langle (\exists v : \phi)(\psi) \rangle\!\rangle^{\mathcal{M}} = \{g[v \to m] \mid g \in \langle\!\langle \phi \rangle\!\rangle^{\mathcal{M}} \cap \langle\!\langle \psi \rangle\!\rangle^{\mathcal{M}}, m \in M\}$

We can say an open formula ϕ is *rendered truth-evaluable* by an assignment g if $g \in \langle\!\langle \phi \rangle\!\rangle^{\mathcal{M}}$. A sentence ϕ (without free variables) is truth-evaluable just in case $\langle\!\langle \phi \rangle\!\rangle^{\mathcal{M}} \neq \emptyset$ (or equivalently $\langle\!\langle \phi \rangle\!\rangle^{\mathcal{M}} = \mathcal{G}$). Note that the definition of the domain of significance of e only appealed to facts about domains of significance, and not about extensions or anti-extensions. It's worth mentioning this is a unique feature the Weak Kleene scheme, in which facts about truth-valuelessness are 'separable' in this way. In other projection schemes, we need to incorporate facts about extensions and anti-extensions in recursively tracking truth-valuelessness, integrating the recursions for domains of significance and satisfaction.

The latter pertinent denotations of formulas relative to a model and assignment pair are mostly given as usual, again with anomalous character pro-

[16]Following convention I use $f[a \to b]$ to denote the function differing from f at most in that $f[a \to b](a) = b$. Also, these definitions again predict that quantified claims are anomalous only when they have only anomalous substitution instances. See n.9 for some reasons we might eventually have to relax this requirement.

jected according to the Weak Kleene scheme. The main exception is in the treatment of quantifiers. In this instance, since we're assuming projection in unrestricted contexts also goes by the Weak Kleene scheme, most work is done by appealing to domains of significance again. Quantifiers are evaluated over domains restricted to exclude substitution instances which would generate truth-valueless status were domains to range more broadly. Given Weak Kleene projection for unrestricted quantifiers, this just means that we should restrict quantifier domains over non-anomalous substitution instances.

Denotations of Formulas

The *denotation of a formula* ϕ *in a model* $\mathcal{M}$ *relative to an assignment* g, written $[\![\phi]\!]^{\mathcal{M},g}$ is U if $\langle\!\langle\phi\rangle\!\rangle^{\mathcal{M}} = \emptyset$, and otherwise is given as follows:

$$[\![R(\tau_1,\ldots,\tau_n)]\!]^{\mathcal{M},g} = \begin{cases} T & \text{if } \langle[\![\tau_1]\!]^{\mathcal{M},g},\ldots,[\![\tau_2]\!]^{\mathcal{M},g}\rangle \in [\![R]\!]^{\mathcal{M},g} \\ F & \text{if } \langle[\![\tau_1]\!]^{\mathcal{M},g},\ldots,[\![\tau_2]\!]^{\mathcal{M},g}\rangle \notin [\![R]\!]^{\mathcal{M},g} \end{cases}$$

$$[\![\neg\phi]\!]^{\mathcal{M},g} = \begin{cases} T & \text{if } [\![\phi]\!]^{\mathcal{M},g} = F \\ F & \text{if } [\![\phi]\!]^{\mathcal{M},g} = T \end{cases}$$

$$[\![\phi\wedge\psi]\!]^{\mathcal{M},g} = \begin{cases} T & \text{if } [\![\phi]\!]^{\mathcal{M},g} = [\![\psi]\!]^{\mathcal{M},g} = T \\ F & \text{if } [\![\phi]\!]^{\mathcal{M},g} \text{ or } [\![\psi]\!]^{\mathcal{M},g} = F \end{cases}$$

$$[\![\phi\vee\psi]\!]^{\mathcal{M},g} = \begin{cases} T & \text{if } [\![\phi]\!]^{\mathcal{M},g} \text{ or } [\![\psi]\!]^{\mathcal{M},g} = T \\ F & [\![\phi]\!]^{\mathcal{M},g} = [\![\psi]\!]^{\mathcal{M},g} = F \end{cases}$$

$$[\![(\forall v:\phi)(\psi)]\!]^{\mathcal{M},g} = \begin{cases} T & \text{if } \{m\in M|g[v\to m]\in\langle\!\langle\phi\rangle\!\rangle^{\mathcal{M}}\cap\langle\!\langle\psi\rangle\!\rangle^{\mathcal{M}} \text{ and } [\![\phi]\!]^{\mathcal{M},g[v\to m]} = T\} \subseteq \\ & \quad \{m\in M|g[v\to m]\in\langle\!\langle\phi\rangle\!\rangle^{\mathcal{M}}\cap\langle\!\langle\psi\rangle\!\rangle^{\mathcal{M}} \text{ and } [\![\psi]\!]^{\mathcal{M},g[v\to m]} = T\} \\ F & \text{if } \{m\in M|g[v\to m]\in\langle\!\langle\phi\rangle\!\rangle^{\mathcal{M}}\cap\langle\!\langle\psi\rangle\!\rangle^{\mathcal{M}} \text{ and } [\![\phi]\!]^{\mathcal{M},g[v\to m]} = T\} \not\subseteq \\ & \quad \{m\in M|g[v\to m]\in\langle\!\langle\phi\rangle\!\rangle^{\mathcal{M}}\cap\langle\!\langle\psi\rangle\!\rangle^{\mathcal{M}} \text{ and } [\![\psi]\!]^{\mathcal{M},g[v\to m]} = T\} \end{cases}$$

$$[\![(\exists v:\phi)(\psi)]\!]^{\mathcal{M},g} = \begin{cases} T & \text{if } \{m\in M|g[v\to m]\in\langle\!\langle\phi\rangle\!\rangle^{\mathcal{M}}\cap\langle\!\langle\psi\rangle\!\rangle^{\mathcal{M}} \text{ and } [\![\phi]\!]^{\mathcal{M},g[v\to m]} = T\}\cap \\ & \quad \{m\in M|g[v\to m]\in\langle\!\langle\phi\rangle\!\rangle^{\mathcal{M}}\cap\langle\!\langle\psi\rangle\!\rangle^{\mathcal{M}} \text{ and } [\![\psi]\!]^{\mathcal{M},g[v\to m]} = T\} \neq \emptyset \\ F & \text{if } \{m\in M|g[v\to m]\in\langle\!\langle\phi\rangle\!\rangle^{\mathcal{M}}\cap\langle\!\langle\psi\rangle\!\rangle^{\mathcal{M}} \text{ and } [\![\phi]\!]^{\mathcal{M},g[v\to m]} = T\}\cap \\ & \quad \{m\in M|g[v\to m]\in\langle\!\langle\phi\rangle\!\rangle^{\mathcal{M}}\cap\langle\!\langle\psi\rangle\!\rangle^{\mathcal{M}} \text{ and } [\![\psi]\!]^{\mathcal{M},g[v\to m]} = T\} = \emptyset \end{cases}$$

Different construals of the projection of truth-valuelessness in unrestricted quantification will generate different, often much more complicated, clauses for enforcing the quantifier domain restriction. For example, if we adopt Strong Kleene projection for quantifiers, domains won't be restricted to elements in an open formula's domain of significance, but to some *superset* of that domain (that is, we will include in the domain of quantification those objects outside the domain of significance which happen not to contribute to anomalous status in unrestricted quantification). The superset in question may shift from quantifier to quantifier. In this case, domains of significance don't produce the restriction in anything like the simple way above, but the information recursively tracked in domains of significance (along with that tracked by denotations more generally) will of course need to be used in characterizing how the domain restriction occurs.

As usual, a sentence ϕ is true (simpliciter) in $\mathcal{M}$, noted $[\![\phi]\!]^{\mathcal{M}} = T$, just in case for all $g \in \mathcal{G}$, $[\![\phi]\!]^{\mathcal{M},g} = T$. Analogously, a sentence ϕ is truth-evaluable (simpliciter) in $\mathcal{M}$ just in case for all $g \in \mathcal{G}$, $[\![\phi]\!]^{\mathcal{M},g} \in \{T, F\}$ (that is, again, if $\langle\!\langle\phi\rangle\!\rangle^{\mathcal{M}} = \mathcal{G}$).

3 Logical Consequence

The domain restriction from anomaly has some interesting implications for issues in philosophical logic. There are two ways the domain restriction may interact with logical form. Regardless of which option one takes, the presence of anomaly generates failures of classical logic in characterizing important classes of natural language inference—failures of a *kind* that no other known phenomenon generates. Additionally, on very weak assumptions about the projection of anomaly in quantified contexts, the interaction of anomaly with quantifiers may supply new motivations for thinking that capturing a useful set of inferences owing to logical form requires the importation of at least some semantic information. Let me take each idea in turn.

In §1, I argued that a speaker's truth-value attributions to (4) and (5) owed to a semantically enforced domain restriction responsive to anomaly.

(4) Bob uprooted everything in his yard and burned it.

(5) Bob burned everything in his yard.

How the truth of (4) and falsity of (5) affect the logical consequence relation depends on how the quantifier domain restriction from anomaly interacts with the logical form of these sentences.

A first construal takes the logical forms of (4) and (5) to be something like (4′) and (5′), as I provisionally assumed in §2.2.

(4′) $[\forall x\text{: InYard}(x, \text{Bob})][\text{Uproot}(\text{Bob}, x) \wedge \text{Burn}(\text{Bob}, x)]$

(5′) $[\forall x\text{: InYard}(x, \text{Bob})][\text{Burn}(\text{Bob}, x)]$

If this is the case, *a straightforward classical inference is violated among truth-evaluable sentences relative to the same context.*[17] Such a failure is arguably uniquely generated by the presence of anomaly. For example, this kind of inference failure is not as clearly manifested by forms of salience-sensitive quantifier domain restriction, context-sensitivity in general, or even by phenomena which

[17]Of course, the logical forms I've given above use generalized quantifiers. But since the point stands if we return to monadic quantification with a conditional, I'll stick to the notation I've used so far in the paper.

otherwise have the potential to motivate a shift to trivalent semantics, such as presupposition failure. Let me say a little more by way of defending this claim.

Context sensitivity may *seem* to provide a wealth of potential failures of classical inferences schemes. For example even the inference from "It's precisely 5 o'clock" to "It's precisely 5 o'clock" may be suspect, since the first utterance can be true and the second false owing to minuscule changes in their contexts of utterance. Similarly, consider the 'inference' from the potentially true utterance of "All the beer is in the fridge" as used before the beer in the basement is made salient to the potentially false utterance of "All the beer is in the fridge" after that beer has been made salient.

Crucially, however, such examples only pose a threat to classical inference schemes on the assumption that context does not make contributions to logical form. If the time of the context of utterance, or the context of utterance itself, forms part of the logical form of any utterance of "It's precisely 5 o'clock", these won't provide counterexamples to classical schemes. They will simply motivate (relatively minor) complications in our conception of logical form. Relative to *fixed* contexts, classical schemes keeping track of common contributions to logical form from context can be preserved. Similar remarks hold for familiar forms of contextual quantifier domain restriction. In fact, Stanley and Szabó have used the binding phenomenon to argue precisely that the influence of context on quantifier domains is mediated through the presence of variables present in logical form.

Other phenomena that *might* motivate a shift to trivalence, such as certain strong forms of presupposition failure, could have more of an effect on classical inference schemes. But the effect of shifting to trivalence alone is not necessarily as damaging to classical logic as one might expect. The complications arising from trivalence often lead to re-defining consequence not as truth-preservation, but as truth-preservation *among truth-evaluable sentences*. This restricted relation largely factors out the influence of a third-truth-value, again enabling us to pick out a quite substantial body of 'inferences' (now reconstrued) which are truth-preserving-among-truth-evaluables in virtue of logical form. The resulting relation, for obvious reasons, tends to vindicate classical inference schemes, fostering the view that classical logic is the logic of truth-evaluables.

By contrast, the truth of (4) and falsity of (5), given the proposed logical forms (4′) and (5′), in some sense constitute as real and substantial a violation of classical logic as one could get. If we adopt (4′) and (5′) and leave the task of effecting domain restriction to the clauses of the semantics for the universal quantifier, we have kept logical form too simple to allow a move like that typically used to safeguard classical schemes in the face of context-sensitivity. Moreover, unlike with other engagements with trivalence, redefining inference to track truth-preservation-among-truth-evaluables won't help in this case, as (4) and (5) are both truth-evaluable. And of course this particular failure isn't the only one—many other failures of standard quantified inferences will have to go by the board increasing the sense that classical inference schemes are inadequate for capturing basic quantified inferences in natural language, even just for "all" and "some".

All this is true if we stick to (4′) and (5′) as the proper logical forms for (4) and (5). But there is a second way of construing their logical forms, one which I didn't explore in §2.2. Rather than enforcing the domain restriction metalinguistically in the semantics of quantifiers, we can take it to be effected by an element realized in the logical form of quantifier restrictors. This would most likely be done by systematically accommodating special variables or functions in quantifier restrictors. The logical form of (4) and (5) might then look something like:

(4″) $[\forall x\text{: InYard}(x, \text{Bob}) \wedge f_i(x)][\text{Uproot}(\text{Bob}, x) \wedge \text{Burn}(\text{Bob}, x)]$

(5″) $[\forall x\text{: InYard}(x, \text{Bob}) \wedge f_j(x)][\text{Burn}(\text{Bob}, x)]$

Given the arguments of §1.2, the values of the variables or functions are *not* here supplied by features of the context of linguistic use, but by the semantics of the sentences themselves. Even so, what is important about this construal is that the apparent classical violation in the blocked inference from (4) to (5) is *merely* apparent: the logical forms of (4) and (5) are more complex than their surface grammar reveals. So, in a way, classical logic is safeguarded.

But we don't merely care about what the logic of our language is, but how often it applies to inferences we actually make. Though the strategy adopted on the second construal avoids *violating* classical inference schemes, it does so without safeguarding its *applicability* to some of the most common natural language inferences. The reason is that on this construal a vast range of quantified natural language inferences have logical forms making them classically invalid. To take just one example, the inference from (23) to (24) is not valid, even holding the contributions of context of use fixed, since their logical forms would be (23′) and (24′) with distinct assignments to f_i and f_j because of the different predicates figuring in each sentence supplying their values.

(23) Every man left.

(24) Every short man left.

(23′) $[\forall x\text{: Man}(x) \wedge f_i(x)][\text{Left}(x)]$

(24′) $[\forall x\text{: Man}(x) \wedge \text{Short}(x) \wedge f_j(x)][\text{Left}(x)]$

Note that the same kind of point doesn't apply to inferences involving sentences with standard forms of context sensitivity, since the inferences that *are* good logical inferences are ones where it's quite plausible that contributions from context to logical form can be held fixed. For example, the same problem won't arise for analogous treatments of contextual quantifier domain restriction, modeled with tacit variables or functions. In those cases, whatever shifts the values of the tacit material (when unbound) is clearly *exhausted* by facts about context of linguistic use. If I safely infer "All the beer is in the fridge or on the counter" from "All the beer is in the fridge", despite a domain restriction in both, we can continue to construe the inference as a logical one since contextual contributions to the restrictions are plausibly the same.[18]

Thus, on this second way of treating the domain restriction, though traditional logic gets the semantics of the quantifiers right, it no longer *on its own* captures anywhere near as substantial and interesting a body of natural language inferences as traditionally conceived. Since the problem stems from the semantic features of the sentences used in inference, we can't sidestep this issue in the way we might for corresponding 'failures' owing to context-sensitivity: it is fruitless to try to preserve a substantial body of inferences by restricting attention to a fixed context. What this means is that on this second proposal, logic as applied to natural language inference would, in effect, be reduced to an awkward and meager extension of propositional logic.

So, regardless of whether the domain restriction from anomaly is built into the recursive clauses for quantifier interpretation, or whether it is mediated by an element in logical form, this domain restriction threatens the applicability of classical logic to natural language inference in ways that no other known phenomenon does.

Appreciating this point should reinforce the idea that anomaly interestingly transforms fairly standard conceptions of logical consequence. What then should a revised consequence relation look like? Our first move in recapturing a set of valid quantified inferences should be relatively straightforward given our work in §2. Letting Γ be a set of sentences and ϕ a sentence, a typical consequence relation $\models$ is given by the following definition:

$$\Gamma \models \phi \Leftrightarrow \text{ for any model } \mathcal{M}, \text{ if } \forall \gamma \in \Gamma \ [\![\gamma]\!]^{\mathcal{M}} = T, \text{ then } [\![\phi]\!]^{\mathcal{M}} = T$$

where $\mathcal{M}$ ranges over bivalent models. To get a better consequence relation, we need only allow $\mathcal{M}$ to range more broadly over the trivalent models, such as those supplied in §2.2, that track the influence of anomaly on the domains of quantifiers.

Such a move, however, brings other problems to light which concern the potential projective behavior of truth-valuelessness. The stronger its projective behavior, the more inference schemes are lost. To take one example, if truth-valuelessness projects across disjunction, the above definition will ensure $\phi \not\models$

[18] Also note that my claims here don't turn on how I'm using the word 'context', namely to apply to general features of context of use. One might grant that general features of the context of utterance needn't supply the values of f_i and f_j, when it comes to the domain restriction, but maintain that features of *linguistic context*—e.g., the words used in each expression—are doing that work. I don't want to contest terminology. That's a perfectly fine way of using the word "context". But relabeling terms doesn't avoid the problem I'm after. When you use "context" in this sense, it's now simply the case that contributions from context *must* change from (23) to (24) (since any transition from (23) to (24) is *ipso facto* one in which contributions from linguistic context are changing). So again, you can never find an inference from (23) to (24) that is underwritten by purely logical relations.

$\phi \vee \psi$. As I've tried to stress, I mostly want to stay neutral on the question of how truth-valuelessness projects in non-quantified contexts, since this is a controversial and delicate matter I don't have the time to discuss. What we can say here is that very minimal assumptions about the projective behavior of anomaly in quantified contexts on their own ensure that anomaly unduly perturbs even a consequence relation redefined over trivalent models as above. Consider:

(25) Few men ate.

(26) * Few men refrained around the discipline and ate.

(27) No buildings are grey.

(28) * No pubescent buildings are grey inversions.

If (26) and (28) (or suitable variants of them) are anomalous, and my arguments from §1 are sound, these examples give violations of quantified inferences which are importantly different from those I've examined so far. The domain restriction from anomaly on its own provides no reason to expect a violation of the inference, for example, from "No Fs are Hs" to "No F-and-Gs are H-and-Is". After all, further winnowing the objects satisfying the restrictor and matrix of "No" typically only increases the likelihood its use will come out true. The problem arises from the fact, noted in §2.1, that quantified claims with only anomalous substitution instances tend to come out anomalous.

As these examples should help reveal, this phenomenon has the potential to greatly perturb the inferential schemes licensed by our definition of consequence above. In general, inferences which move between quantifiers while adding more material to the quantifier scope or restrictor will be threatened. This results in a set of valid quantified inferential schemes that looks erratic, and somewhat uninteresting—perhaps even more erratic and uninteresting than the set of valid propositional inferential schemes due to the introduction of a third truth-value.

To cope with this problem, it is natural to take a strategy I alluded to earlier: since truth-valueless whole sentences tend to render the consequence relation uninteresting, we can get a better grip on the class of inferences mediated by logical form by simply 'factoring out' the influence of these problematic sentences on the consequence relation. That is to say, we can recharacterize the consequence relation as one which relates *truth-evaluables*, as follows.

$$\Gamma \models_f \phi \Leftrightarrow \forall \mathcal{M},\ \text{if } \forall \gamma \in \Gamma\ [\![\gamma]\!]^{\mathcal{M}} = T \text{ and } [\![\phi]\!]^{\mathcal{M}} \neq U,\ \text{then } [\![\phi]\!]^{\mathcal{M}} = T$$

I'll call $\models_f$ the relation of *formal logical consequence*, for reasons that will be clear soon.

Formal consequence succeeds in picking out a significant class of schemes conducive to inference including many involving quantified sentences, while blocking those from the domain restriction due to anomaly. For example, as long as anomaly projects over conjunction, and unrestricted universal quantifiers inherit truth-valuelessness from any truth-valueless substitution instances in their matrix,

$$[\forall x : Fx][Gx \wedge Hx] \not\models_f [\forall x : Fx][Gx]$$

in line with the evaluations of (4) and (5) in §1. On the other hand, other quantified inferences such as

$$[\forall x : Fx][Gx] \models_f [\forall x : Fx \wedge Hx][Gx]$$

are safeguarded.

What the final relation looks like, of course, depends on our choice of projection schemes. As an example, and to follow through on the scheme I've been working with so far, if we adopt a Weak Kleene scheme for propositional connectives, the formal consequence relation only eliminates quantified validities from the bivalent setting that need to be jettisoned due to the quantifier domain restriction from anomaly.

Proposition 3.1. *Let* BIVALENT *denote the set of valid bivalent inferences,* PROP *denote the set of valid inferences of bivalent propositional logic, and* FORMAL *denote the set of formally valid inferences (for trivalent models using a Weak Kleene scheme as in §2.2). Then the following relations hold:*

$$\text{PROP} \subsetneq \text{FORMAL} \subsetneq \text{BIVALENT}$$

Proof. Suppose Γ propositionally entails ϕ in the bivalent setting, and a trivalent model $\mathcal{M}$, of the kind given in §2.2, is such that $\forall \gamma \in \Gamma, [\![\gamma]\!]^{\mathcal{M}} = T$, and $[\![\phi]\!]^{\mathcal{M}} \neq U$. Then, since we are working in a Weak Kleene scheme, each truth functional component θ_i of ϕ or formulas of Γ is such that $[\![\theta_i]\!]^{\mathcal{M}} \neq U$. But then we are essentially in the bivalent case, so we have $[\![\phi]\!]^{\mathcal{M}} = T$. This shows the first containment. The second containment follows from the fact that bivalent models are just trivalent models of §2.2 with degenerately broad domains of significance. That the containments are proper is witnessed by the two examples of which quantified inferences are, and are not, formally valid given just above. □

Other schemes may of course result in a very different formal consequence relation.

Adopting the formal logical consequence relation, however, comes with an important philosophical cost. $\models_f$ does *not* model sound inference, but only sound inference *among truth-evaluables.* So there will be very many sentences ϕ and ψ, and models $\mathcal{M}$, such that $\phi \models_f \psi$ while $[\![\phi]\!]^{\mathcal{M}} = T$ and $[\![\psi]\!]^{\mathcal{M}} \neq T$. This is significant because of a standard construal of what a logical consequence relation *should be.*

Formal logic, in the sense I'm alluding to, is conceived as in the business of tracking which inferences are truth-preserving in virtue of logical form. It does this through attention to how the truth-conditions of complex sentences systematically correlate with aspects of their logical form. Logic, though sometimes employed in conjunction with semantics to swell the relevant body of inferences tracked, is thought to have an autonomous domain. The idea, from the ancients down to the early analytics, is that there is a substantial and interesting body of inferences which are *entirely* content or subject-matter independent. The idea that there is such a body of inferences is potentially threatened in unique ways by the projection behavior of anomaly. The influence of anomaly in generating truth-valueless quantified claims threatens to make a class of quantified inferences which appeals *solely* to logical form look impoverished and erratic, as already noted.

What makes the set uninteresting is a problematic interaction between two desiderata: aiming to track pure logical *form* conducive to inference, and aiming to track a class of truth-preserving inferences *unto themselves*. To adopt the formal consequence relation is to concede the force of this tension, and jettison the second of these desiderata in favor of the first. We can recapture genuine truth-preserving inference with a second consequence relation, which I'll call the *semantic logical consequence* relation, that imports a minimal amount of semantic information to capture genuine truth-preserving inference due to logical form.

$$\Gamma \models_s^{\mathcal{M}} \phi \Leftrightarrow \Gamma \models_f \phi \text{ and } [\![\phi]\!]^{\mathcal{M}} \neq U$$

By appealing to a model parameter, we can capture information about truth-evaluability needed to ensure that a transition from Γ to ϕ is one which is guaranteed to preserve truth, and by appealing to $\models_f$ we ensure the transition is indeed mediated by logical form, to the extent logical form can make contributions to inference. Just as with $\models_f$ though, $\models_s^{\mathcal{M}}$ sacrifices one intuitive hallmark of a logical consequence relation for another: It captures all and only sentence transitions which are genuinely truth-preserving due to their logical form. However, it does this at the expense of incorporating semantic information via the model parameter, and thus can no longer be construed as tracking inferences which are *entirely subject matter independent*.

The fact that the semantic consequence relation imports information from a model has important philosophical consequences. For it shows that assessing whether a legitimate inference has been drawn between sentences *may require basic information about the semantics of those sentences*. Put another way: we cannot simply look at the syntactic features of sentences to discover information about truth-preserving inference.

I don't have the space to discuss the full importance of these issues here. I have merely wished to call attention to the fact that anomaly's interaction with quantification has two potentially interesting implications for philosophical logic. On the one hand, the ways in which anomaly *does not* project in quantification afford us our most substantial threat to the utility of classical inference schemes. On the other hand, the ways in which anomaly *does* project in quantification may provide us with special reasons to doubt that there are substantial and regular bodies of sentence to sentence transitions which preserve truth *solely* in virtue of their logical form.

References

K. Bach (1994). 'Conversational Impliciture'. *Mind & Language* **9**(2):124–162.

K. Bach (2005). 'Context *ex Machina*'. In *Semantics vs. Pragmatics*, pp. 15–44. Clarendon Press, Oxford.

D. Beaver (2001). *Presupposition and Assertion in Dynamic Semantics.* CSLI Publications.

E. Camp (2004). 'The Generality Constraint and Categorial Restrictions'. *Philosophical Quarterly* **54**(215):209–231.

R. Carnap (1937). *The Logical Syntax of Language.* London: Routledge & Kegan Paul.

G. Chierchia & S. McConnell-Ginet (2000). *Meaning and Grammar.* MIT Press, Cambridge, MA, 2nd edn.

J. Fodor & J. Katz (1964). 'The Structure of a Semantic Theory'. In J. Fodor & J. Katz (eds.), *The Structure of Language: Readings in the Philosophy of Language*. Prentice-Hall, Englewood Cliffs, NJ.

G. Frege (1892). 'On Sense and Reference'. *Translations from the Philosophical Writings of Gottlob Frege* **3**.

G. Frege (1918). 'The Thought: A Logical Inquiry'. *Mind* **65**(259):289–311. reprinted 1956.

M. Glanzberg (2006). 'Quantifiers'. *The Oxford Handbook of Philosophy of Language* pp. 794–821.

H. Grice (1991). *Studies in the way of words.* Harvard University Press Cambridge, Mass.

I. Heim & A. Kratzer (2004). *Semantics in Generative Grammar.* Blackwell, Malden, MA.

L. Horn (1985). 'Metalinguistic Negation and Pragmatic Ambiguity'. *Language* **61**(1):121–174.

S. Lappin (1981). *Sorts, Ontology, and Metpahor: the Semantics of Sortal Structure*. Walter De Gruyter.

W. V. O. Quine (1986). *Philosophy of Logic.* Harvard University Press.

R. Routley (1966). 'On a Significance Theory'. *Australasian Journal of Philosophy* **44**(2):172–209.

B. Russell (1905). 'On denoting'. *Mind* **14**(56):479–493.

B. Russell (1908). 'Mathematical logic as based on the theory of types'. *American Journal of Mathematics* **30**(3):222–262.

G. Ryle (1984). *The Concept of Mind*. University of Chicago Press, Chicago, IL.

J. Stanley & Z. G. Szabó (2000a). 'On Quantifier Domain Restriction'. *Mind and Language* **15**(2 and 3):219–261.

J. Stanley & Z. G. Szabó (2000b). 'Reply to Bach and Neale'. *Mind & Language* **15**(2&3):295–298.

P. F. Strawson (1950). 'On Referring'. *Mind* **59**(235).

R. Thomason (1972). 'A Semantic Theory of Sortal Incorrectness'. *Journal of Philosophical Logic* **1**(2):209–258.

K. von Fintel (1994). *Restrictions on Quantifier Domains*. Ph.D. thesis, University of Massachusetts.