

What is a Truth-Value Gap?

James R. Shaw

Draft, please don't cite

Truth-value gaps are generally taken to be a well-understood phenomenon. From a formal perspective this is correct. Gaps are truth-values that block inferences as falsehood does while having a stronger projection behavior.

But from a foundational perspective gaps are actually very poorly understood. What exactly is the *point* of distinguishing between more than one way of being untrue? How do gaps fit into, and arise out of, our theories of assertion or the nature of content? Difficulties in supplying answers to such questions have given rise to several challenges that philosophically interesting truth value gaps *can't* exist. My goal in this paper, beginning with one of the most powerful challenges of this kind by Dummett, is to offer one account truth-value gaps that begins to address some of the more pressing foundational questions about them. The account yields some important conclusions about what resources we need in our theories of assertion to posit gaps, the relation they have to mentality, and the sources of their compositional behavior.

1 Dummett's Challenge

In “Truth”, Dummett bemoans attempts to explain truth by specifying which statements are true under various conditions.¹ After all, this would simply classify statements without yet giving a purpose to that classification.

To highlight the problem, Dummett draws an analogy with competitive games. If a theorist of such games divided final states of play into three classes—*win*, *lose*, and *draw*—that classification itself wouldn't go far to illuminate competitive games, or the various states themselves. Any temptation to think so comes from tacitly drawing on an understanding of its labels. If instead the theorist had labeled the states of play the *A*, *B*, and *C-states*, it would become clear what part of the explanation of the classification had gone lacking: that prototypically players *aim* at the *A-states*. If instead the *B* or *C-states* played that role, we would have a very different class of competitive games on our hands. And if no states served that function, we wouldn't have a classification pertinent to competitive games at all.

¹Dummett (1978), whose discussion I'll follow closely here. Dummett's argument has met with much approval, especially in the literature on the semantic paradoxes. See, for example, Glanzberg (2003) and Priest (2006) both of whom offer sharpened versions. For now, I'll provisionally follow Dummett's use of “statement” to pick out bearers of truth.

In stressing this, Dummett calls attention to the fact that we haven't explained truth, falsehood, or any other truth-value until we have specified the point of applying them. There may be several ways of labeling statements that yield different and perhaps overlapping classifications. So if we are going to draw philosophical conclusions from some allotment of truth-values, we must know what the allotment is tracking. A key reason Dummett harps on this idea is because he thinks that our purposes in allotting truth values force a binary structure on them. This means that gap theorists positing a third truth-value are guilty of manipulating an empty formalism or of conflating divergent classificatory schemes.

Dummett claims the purpose of assigning truth-values is to mark a role in the determination of assertoric content (for Dummett, sense), and that without providing the details of such an account, we can see the operations of assertion leave no room for the existence of truth-value gaps. In particular, Dummett claims:

A statement, so long as it is not ambiguous or vague, divides all possible states of affairs into just *two* classes. For a given state of affairs, either the statement is used in such a way that a man who asserted it but envisaged that state of affairs as a possibility would be held to have spoken misleadingly, or the assertion of the statement would not be taken as expressing the speaker's exclusion of that possibility. If a state of affairs of the first kind obtains, the statement is false; if all actual states of affairs are of the second kind, it is true. It is thus *prima facie* senseless to say of any statement that in such-and-such a state of affairs it would be neither true nor false.²

The idea seems to be that truth values record when a conventional act rules in, or rules out, various possibilities. Modifying this idea provisionally to suit the idiom of propositions and possible worlds:

- (t_d) *A*'s statement that *p* is true-at-a-world *w* if it would not conventionally express *A*'s exclusion of *w* as a possibility.
- (f_d) *A*'s statement that *p* is false-at-a-world *w* if it conventionally expresses *A*'s exclusion of *w* as a possibility.

These conditions would yield a neat argument for the exhaustiveness of truth and falsity conditions in classifying assertoric content. After all, the requirement on falsity is just the negation of the requirement for truth. Exhaustivity would require that the excluded middle to apply to the requirement for falsity-at-a-world, but this is not implausible. And Dummett rightly points out that this characterization already does substantial work ruling out several phenomena alleged to require gappy treatment. For example, statements with empty

²Dummett (1978) p.9.

names on the above criteria are arguably false (one speaks misleadingly in using name without excluding cases where the name doesn't refer), and conditionals with false antecedents are clearly true (speakers asserting "if p , then q " doesn't thereby rule out that p is false).

For now, I want to set aside the question of whether (t_a) and (f_a) are faithful to Dummett, and even the question of whether the conditions are good ones. This is because the brunt of Dummett's argument doesn't turn on the precise formulations of such conditions. Rather, his core point is this: if we use truth-value classifications to track a single kind of assertoric status or effect, which does not come in degrees, then those classifications will be bipartite. Dummett thinks that reflection on the point of assertion reveals that the philosophically interesting uses of truth values track a single kind of assertoric effect—something like ruling out alternatives. As such, going beyond a bipartite classification of assertoric effects becomes superfluous.

If the goal or consequences of an act are varied, or coming in degrees, trivalent classifications regain their utility. Dummett thinks that this occurs for conditional bets, like "If John comes in the next few minutes, I bet it will be without Sally." Tracking the effects of conditional bets requires a tripartite scheme because they can have three distinct outcomes: payment, receipt of payment, and no exchange. By contrast: "Statements are... not like bets; the making of a statement has, as it were, only one kind of consequence."³ The consequence is, as I've said, something like ruling out possibilities. And the uniqueness of this consequence precludes any significant tripartite classification of assertoric effects.

Dummett's argument, though simple, is incredibly powerful. To appreciate its force, we can note how it immediately blocks a standard form of argument in favor of gaps from projection phenomena. Many have held that we need truth-value gaps to explain the embedding behavior of sentences, especially under negation. Negation seems to transform the uncontroversially true into the uncontroversially false, and vice versa. But some statements seem to exhibit a kind of imperfection that is preserved under negation—those involving empty names, for example. If truth-values are used to track compositional effects in assertion, it is sometimes claimed, we will need gaps in addition to truth and falsity to model a third way in which whole sentences interact with negation.

Dummett notes that, provided his argument is sound, even if we do need some third status for *compositional* purposes we will end up assimilating that third status into a *species* of truth or falsehood at the level of assertoric content. To appreciate the point, consider an artificial case in which we stipulate the existence of a species of falsehood with a relevant kind of aberrant projective behavior. Consider a fictitious word "squeen", with truth-values assigned to its uses as follows, with G symbolizing "is green" and S "is square".

	$Ga \wedge Sa$	$Ga \wedge \neg Sa$	$\neg Ga \wedge Sa$	$\neg Ga \wedge \neg Sa$
a is squreen	t	t	f^*	f

³Dummett (1978) p.12.

Let's stipulate that whether an utterance is assigned t or the values in $\{f, f^*\}$ is tracking *effects in assertion*. So there is no difference between saying “ a is squeen” and “ a is green”. But the division between f and f^* marks a difference in interaction with a negation-like operator, pronounced “neg”, so that we have the following:

	$Ga \wedge Sa$	$Ga \wedge \neg Sa$	$\neg Ga \wedge Sa$	$\neg Ga \wedge \neg Sa$
a is neg squeen	f	f	f^*	t

So there is no difference between saying “ a is neg squeen” and “ a is neither green nor square”.

Whether we use words like “neg” and “squeen” is clearly an empirical question for compositional semanticists. If we do, we must go beyond a bipartite assignment of extensional values to sentences for compositional purposes. But this *by itself* won't settle any questions about assertion. In the squeen example, only bipartite classifications are needed to model the effects of assertions containing “squeen”: every sentence containing that word is assertorically equivalent to a sentence having straightforwardly two-fold effects *by stipulation*.

Now Dummett's point is this: his arguments, if successful, establish that no matter how complex the compositional behavior of whole sentences, once we consider the effects of a composite expression in assertion *every* case must be like the “squeen” case. We might need four, eight, or continuum many values in compositional semantics to explain various forms of embedding, but effects in assertion will always be a ‘yes-no’ matter.

But in this case we may have two distinct classifications: a ternary one for compositional values, and a binary one for assertoric values. If the classifications come apart, which are the ‘genuine’ truth values? Here our methods recommend tolerance of alternatives: use the words and corresponding classifications however you want, as long as you are clear about their use. But Dummett insists, and here I agree, that for the most part the existence of gaps will only do the work *philosophers* have historically demanded of them if they are needed to characterize assertoric and not compositional effects. Consider, for example, a paradigmatic philosophical topic to which gaps are applied: the semantic paradoxes. If gaps only arise for compositional purposes, they are useless in helping to diagnose paradox, or model the existence of borderline cases, since the problems in neither case stem exclusively from compositional concerns.

If Dummett is right that philosophically important truth-value classifications mark off roles in the determination of assertoric effects and there is, as Dummett puts it, just one *kind* of consequence to an assertion, which does not come in degrees, then philosophically interesting truth-value gaps can't exist, and nothing can count as evidence for them. This is because no sense can be made of the phenomenon that the evidence is alleged to be evidence *for*.

2 Grades of Assertoric Effect

For expositional purposes, I've followed Dummett in treating statements as truth-bearers, which is a debatable assumption. There are at least three possible candidates worthy of consideration: utterances of sentences, speech acts in which they are produced, or 'what is expressed' by such a speech act—typically something more abstract, like a proposition. I want to remain agnostic for now as to which of these we should choose. A main reason for this is the key lesson of the previous section: that we shouldn't be overly attached to associating words like "true" with a single classificatory role—the importance of any classification stems from its purpose, not the pre-theoretic labels we use for it. What I eventually want to do on the gap-theorist's behalf is to carve out some tripartite classification relevant to the operation of assertion that will do the work some philosophers have traditionally wanted it to. We can give the gap theorist more latitude, and avoid *some* controversies, by provisionally remaining agnostic as to what is to be classified.

Since the primary desideratum for the gap-theorist's tripartite scheme is that it be relevant to assertoric practice and not (merely) compositional role, I'll provisionally attribute truth values to utterances used in assertion, and talk of what is thereby tracked as an *assertoric effect*, this being a blanket term to cover something that happens due to the assertion or a status the assertion has.

Having clarified my terminology, I now want to argue that Dummett's challenge poses even more problems for gap theorists than is usually acknowledged. To help bring out why, let's begin by formulating Dummett's argument with the following two key premises.

- (A) If the point of philosophically significant allotments of truth-values is to track the presence or absence of a single kind of assertoric effect that doesn't come in degrees, then tripartite classifications of truth-values are theoretically unmotivated.
- (B) The point of philosophically significant allotments of truth-values is to track the presence or absence of a single kind of assertoric effect that doesn't come in degrees.

(A) is an unexceptionable claim, though one whose importance is easy to overlook. In voicing it, Dummett is rightly reminding gap theorists what kind of account they owe us: one which ties the structure they impose on distributions of truth-values to aspects of the assertoric speech acts those truth-values allegedly model.

(B) is more controversial, and will be the premise that gap theorists must resist. They can't deny that truth-values model assertoric effects. This would threaten to rob gaps of their most important philosophical applications, as I've already noted. This leaves only two possible ways for gap theorists to resist Dummett's devastating conclusion.

- (I) Maintain that truth-values model a single assertoric effect that comes in degrees.

(II) Maintain that truth-values model multiple distinct assertoric effects.

(I) probably looks like the safest route to make space for gaps. The next thing I'd like to do is show, surprisingly, that appealing to (I) should be avoided by gap theorists if at all possible.

The reason why (I) should look initially tempting is because of the most common informal characterization of truth-value gaps in the literature: as involving some kind of *indeterminacy*. Virtually any effect can exhibit the indeterminacy in question. My command might get you to run at least 10 miles or not. But what if the surface you're running on is bumpy? What speed must you maintain to have *run* the distance? Questions like these make it look as if there are cases where my words clearly had the effect in question, cases where they clearly did not, and 'intermediate' cases which are hard to classify either way.

From this perspective every effect comes in at least three (if not more) degrees. And we can apply this idea more directly to assertoric effects. Suppose that assertions have some unique characteristic function—say, to rule possibilities out. Then we might not only need to track when an assertion conventionally rules some possibility out, or does not, but also 'intermediate' cases where an assertion only indeterminately has the effect in question.

Do such indeterminacies in assertoric effects exist, and are they worth modeling? Probably. But should gaps be *identified* with these indeterminacies? Absolutely not. The problem is that the indeterminacies are not present in the grand majority of cases that gap theorists do, and should, want to model. This might seem surprising. Aren't gaps paradigmatically used in theories of vagueness? And don't vague uses of language clearly involve the indeterminacies of which I've been speaking?

The answer both of these questions is "yes". The problem is that what leads us to press them is the conflation of two distinct kinds of indeterminacies: indeterminacies in assertoric effect, and indeterminacies in *linguistic use*. The simplest and clearest cases of vagueness involve the latter kind of indeterminacy. Imagine, for example, that we have as clear a case as possible of someone who is a borderline case of baldness—Chuck. Speakers are wholly, and uniformly, reluctant to talk about Chuck using the word "bald". They refuse to call him "bald". They refuse to call him "not bald". They insist that he's an intermediate case—dead center between being clearly bald and clearly not. They even insist that because of this "bald" is simply not the right word to talk about how things stand with Chuck.

If this kind of case were to arise, it would be a clear manifestation of the vagueness of "bald", and a phenomenon that gap theorists would be interested to capture. The problem is that there is no reason to think that calling Chuck bald has *indeterminate assertoric effects* (at least at the actual world). To see this, consider the suggestion I adapted from Dummett that falsity-at-a-world w is used to model when a statement is used to conventionally rule w out. If we take this to be the sole effect that truth-values are used to track, then if someone were to assert "Chuck is bald" (by assumption, probably someone who hasn't recently seen him), they would clearly not be performing an act

which conventionally ruled out the actual world. Better yet, to use Dummett's own words, is it the case that someone "who asserted ["Chuck is bald"] but envisaged... the [actual] state of affairs as a possibility would be held to have spoken misleadingly"? Clearly, almost by the very stipulations of the case. So by either Dummett's proposed classification of effects, or my reworking of those effects, Chuck's state does not present us with a case of indeterminacy of assertoric effects.

The problem doesn't just arise from those particular choices. Any clear way of classifying assertoric effects tends to face the same problem. For example, one might think that truth-values are tracking not what speakers are ruling out, but what speakers should aim at in producing their assertions (an idea I'll return to shortly). One might think, for example, that truth-at-a-world w tracks the kinds of assertions we should aim to produce were w actual. Again if this is the only effect we are tracking, assertions of "Chuck is bald" clearly lack this status with respect to the actual world, and so again will not produce truth-value gaps if we identify them with indeterminacies of the single proposed assertoric effect.

I want to clarify that I am not claiming that indeterminacies of assertoric effect should not be modeled with truth-value gaps. I am simply claiming that the converse doesn't hold: the grand majority of gaps do not result from indeterminacies in a single assertoric effect, and so we should not identify gaps with those indeterminacies. The problem here is not limited to cases of vagueness, which are not the only ones we would hope to salvage on the gap theorist's behalf. Difficulties arise in a similar way for conditionals with false antecedents, uses of empty names, presupposition failure, liar-like paradoxes, and category mistakes.

We already discussed this point earlier: someone using a conditional with a false antecedent is not making an assertion that conventionally rules out the worlds at which the antecedent is false. Nor (to try out our second proposed effect) should we be striving to produce them in circumstances which make the antecedents false. Once we fix the single kind of effect truth-values will be used to track, the grand majority of cases gap theorists want to treat will not generate gaps at all because we can comfortably classify them as having, or lacking, the effect in question.

The appeal to indeterminate effects results from a failure to fully appreciate this point, and relies too directly on special features of vague terms as a source of gaps. It should have been clear that this might be problematic for gap theorists, as the majority of the cases to which gap theories have been applied can arise in languages which are completely precise. If we use gaps only to model indeterminacies in assertoric effects, all of these applications of gap theories have to go by the board.

The simplest way to meet Dummett's challenge is to agree that assertions have a single effect, but maintain that tripartite classifications could still be theoretically significant because all effects can be indeterminate, and hence come in at least three degrees. I've argued that this defense, even if successful, would be a Pyrrhic victory for gap theorists, leaving them with almost no viable applications of their view. But aren't there other ways to appeal to grades of

assertoric effect? Why not simply appeal to an effect that more naturally comes in degrees having nothing to do with indeterminacies? For example, appealing to effects tied to probabilistic information might have this kind of structure.

The problem with this route should be even more apparent. What gap theorists want is a tripartite division of assertoric effects, where the third grouping of effects is *defective* in some way and exhibits special projective behavior. Assertoric effects that are more naturally graded have neither of these features. Each grade is equally ‘standard’ for assertoric purposes, and also tends to lack gap-like projection behavior. So there is no reason to think that more natural gradings of assertoric effects will be of any help here.

Gap theories won’t survive in any recognizable form if gaps are construed as grades of a single assertoric effect. That strikingly leaves gap theorists with only one option to meet Dummett’s challenge: to proliferate assertoric effects and show how their combinations give rise to a status with the defective and projective behavior they want. This conclusion is surprising since, to my knowledge, no gap theorist has so much as considered thinking of gaps in these terms. Perhaps even more surprising, though, is how natural and appealing a theory results from taking this option.

3 Amalgamating Assertoric Effects

Dummett claims that making a statement has a ‘single kind of consequence’. I think this is incorrect. We typically have a *compound* goal in making assertions, which generates a corresponding amalgamation of two interacting effects. In special cases, we can pry these two effects apart, thereby meeting Dummett’s challenge and uncovering an important phenomenon for a third truth-value to represent. In this section I’ll describe the tripartite classification I’m alluding to. I’ll defer questions about why this classification merits labels like “true” and “false”, to §5.

Until now, because I have been arguing *with* Dummett, I’ve been able to remain agnostic about what exactly assertoric effects are. We’ve seen that without being committal about the nature and function of assertion, we can dramatically constrain gap theorists’ options. In giving a positive story on behalf of the gap theorist, though, I’ll have to start taking a more controversial stand on how assertion works. To do this, rather than engaging with Dummett’s account, I’d like to build my positive case by drawing on an attenuated version of a familiar picture of the mechanics of assertion developed by Stalnaker.⁴

On that picture an assertion is a kind of rational action which takes place against a shared background of information—information that all parties are commonly aware is being taken for granted for the purposes of conversation. In the paradigmatic case, conversational participants are investigating the world, and how things stand in it. The shared background of information in conversation is a measure of how far the participants are, collectively, in ascertaining just

⁴See, e.g., Stalnaker (1978, 2002).

how the world is. The pool thus determines a set of possibilities—the set of possible ways the world could be that are compatible with the information pooled so far. An assertion is a proposal to augment the shared pool of background information, by conveying information about the open possibilities through a conventional use of language. To have its characteristic effect, the assertion should conventionally do at least two things: partition the open possibilities, and induce a kind of ‘polarity’ on that partition. The polarity marks which side of the partition is claimed to be incompatible with the way the world actually is. The assertion is, at least in part, a proposal to rule those possibilities out.⁵

Already within this simplest outline the framework we can see that prototypical assertion aims at the transfer of *accurate information*—where accuracy and information can to some extent be thought of as independent features of good assertion, or independent goals of the act. On the one hand, I am not doing what is prototypically involved in assertion if I am simply speaking accurately. If that were the only point of assertion, we might all go around telling each other repeatedly that two and two make four. Our goal in asserting is to speak in some appropriate way while providing information, or news, to our interlocutors. Of course what counts as news changes from conversation to conversation. But what remains constant, at least in paradigmatic cases, is that information involves *ruling out* alternatives. So accuracy isn’t all that matters.

On the other hand, prototypical assertion involves more than *simply* blurting out news, no matter how informative. If I told you the universe has, is, and always will consist entirely of a single pea, I would be giving you a fantastic amount of information—allowing you to rule out all but a small class of related possibilities. Of course, this wealth of information would be irrelevant to how things actually stand, and to tell you this would be highly uncooperative. In fact, you can’t even understand someone as *providing* information in the way characteristic of assertion unless you see them as at least pretending to strive for accuracy.

These two goals—information potential and accuracy—are both so integral to the purpose and function of assertion for us that it *could* be argued that both are constitutive aims of it, at least when considered as a practice. If we only had appropriateness constraints for assertion, there would be little to distinguish asserting from, say, greeting. If we only had informativeness constraints for assertion, there would be little to distinguish it from imaginative play. I don’t want to argue for the constitutivity claims here, since they’re not necessary to make my ensuing points. I’m simply bringing up the issue to flag how very important these two features are to assertion’s characteristic functions for us.⁶

⁵It is worth noting, as I mentioned, that so far this view of assertion is an *attenuated* version of Stalnaker’s own. It is perfectly compatible with my characterization of assertion here that it do *more* than partition possibilities as Russellians and Fregeans might maintain.

⁶This marks an important distinction between my methodology and that adopted by Glanzberg (2003), who claims that we should strive to make sense of truth-value attributions in terms of the ‘intrinsic purposes’ of assertion. I’m agnostic as to whether my two features of assertion meet Glanzberg’s strong standard. This does not deter me because the importance of my classificatory scheme follows directly from the importance of the effects I’m tracking—whether or not they involve intrinsic purposes. As I’ll argue in §5, the importance

Now, on this view of assertion, what *could* truth-values be used track? The route I'd like to pursue is to take truth-at-a-world to track accuracy of the utterance made in assertion—the information it supplies counts as good information, and the act of supplying it rightly performed, when the actual world is among the worlds labeled true. Likewise falsity-at-a-world marks the places where an utterance has informational clout—the ‘ruling out’ effect that makes news *news*. These two effects are doubtless worth classifying in their own right, but are they exhaustive?

The answer turns on whether we can ‘pry apart’ the composite purpose of assertion to yield a finer grained structure. I think in one way we can, but in another we cannot. Call a world w “accuracy inducing” for an asserted utterance if the assertion would be accurately performed were w actual.⁷ Call a world w “informationally efficacious” for an asserted utterance, if making the producing the utterance involves making a conventional proposal to rule out w . There is no such thing as being accuracy inducing and informationally efficacious at the same time: the purpose of ruling out a world precludes the information from being accurate were w actual. But it *does* make sense for a world to be non-accuracy-inducing and non-informationally-efficacious. And I think that this is what truth-value gaps can be used to track. So my proposal, yet to be defended, is that the composite effects of an assertion can require a tripartite classification as follows:

- (t) An utterance U is true-at-a-world w if the information supplied by U would constitute accurate information were w actual.
- (f) An utterance U is false-at-a-world w if the information supplied by U would not constitute accurate information were w actual, but U still *provides information* about w —it is part of a proposal to rule w out.
- (u) An utterance U is defective-at-a-world w if the information supplied by U would not constitute accurate information were w actual and further U is *informationally inert* at w —it isn't part of a proposal to rule w out.

On this view what distinguishes falsity and defectiveness is informational potential, and what distinguishes truth and defectiveness is their contribution to accuracy assessment. This is why defectiveness-at-a-world cannot be assimilated to the effects tracked by either of the other two truth-values. But interestingly defectiveness *shares* a feature with each of the standard truth-values: inaccuracy-at-the-actual-world with falsehood, and informational-inertness-at-a-world with truth. I'll return to these latter ideas shortly.

With the proposal outlined, let me motivate the existence of this tripartite division of assertoric effects with an analogy involving a kind of cooperative

of the effects makes the classificatory scheme of sufficient interest to semanticists, logicians, and philosophers.

⁷A useful way to test this is by asking: would it be correct or appropriate to make the assertion if w were *known* to be actual? We needn't pronounce further on what accuracy comes to, beyond what is needed to understand its role in the picture of assertion I've already outlined.

endeavor with a structure very much like assertion. The point of the example is *just* to show how it can make sense to pry apart effects of accuracy and information in the way suggested by (t), (f), and (u).

Consider a game involving two players, *A* and *B*, who are confronted with an array of marbles one of which is known to be privileged by *A*, but not *B*. *A*'s job is to convey to *B* which marble is privileged using a series of moves, involving the creation of a pile of marbles to the left and a pile of marbles to the right. The idea is that the privileged marble must be contained in the pile to the left, and not to the right.

For now I'll suppose the sorting strategy isn't governed by rules in the game—the players just happen to be mutually aware, for whatever reason, that this is how best to convey to *B* where the privileged marble is given other strictures of the game. And we can suppose those strictures to include rules somehow precluding *A* from placing the privileged marble alone on the left.

By convention when *A* divides *all* the marbles in two piles, she performs an act with informational clout by ruling out the the pile on the right as containing the privileged marble. This exhausts the information *A*'s action conventionally supplies about the marbles for *B*. But *A* has also performed an act which has something like accuracy conditions: her act was a good one, given *A* and *B*'s mutual understanding, so long as the privileged marble is among those on the left.

Now in the sorting process *A* might for some reason—laziness, confusion, forgetfulness, whimsy, or inability—divide only *some* of the marbles into the two piles, leaving the others scattered in their original positions. In this circumstance we can ask a pair of questions about each marble in the state of play. First: is that marble contributing to the information *B* obtains from the act? Second: would the placement of the marble be accurate were it the privileged one?

Just as in a standard complete sorting, the marbles in the right pile of the partial sorting are responsible for the informational clout of the action. From the act in question the only *news* *B* can gather is that the privileged marble is not in that pile. From this perspective the unsorted marbles closely resemble the pile sorted on the left: marbles in neither pile contribute to the information-potential of the act.

But also just as in a standard complete sorting, marbles on the left are responsible for the relevant kind of accuracy assessment of the action. If the privileged marble is not among those in this leftmost pile, then by the conventions of the sorting procedure, it has not been appropriately placed. This is equally true if the privileged marble is among the unsorted marbles, as the shared aims of the players and their expectations *require* the privileged marble on the left. The privileged marble left unsorted is improperly treated, and to this extent the marbles in the middle more closely resemble the marbles on the *right*: the privileged marble must be excluded from both piles for its placement to be *accurate*.

From the standpoint of tracking its intended effects, the partial sorting move of player *A* into *left*, *right*, and *other* in this game is thus not equivalent to any full sorting move into *left* and *right* only. Assimilating the *other* marbles to the

right marbles ensures accuracy conditions on the placement of the privileged marble are preserved, but changes the sorting's informational effects. Assimilating the *other* marbles to the *left* marbles distorts accuracy conditions on the reallocated marbles, but keeps informational effects constant. The various partial sorting moves should not be taken as equivalent to each other, at least from the theorists' point of view. Though in some sense equally imperfect as departures from the conventions of standard sorting, they nonetheless differ in terms of their informational effects, and as to whether they involve accurate placement of the privileged marble. The game thus represents a case of assertion-like endeavor which can exhibit the tripartite structure that comes from dissociating the two elements in the compound effect of assertion.

Though the game does this, it is worth dispelling a few misunderstandings about how it works. First, the use of two piles in the example is primarily to make salient how my third status can arise. It is irrelevant that we can describe a similar case of accurate information transfer about the marbles so that the actions of the players involve only one pile for sorting (say the 'ruling out' pile only). This is in part because the kinds of assertoric imperfections that the two-pile case makes salient can just as easily arise in a one-pile case. The easiest way to see this is by returning to consider indeterminacies.

Consider, for example, the question of how close to the single pile a marble must be moved to count as in that pile. Marbles whose status in being moved is ambiguous for both *A* and *B* will play the exact same role in their proceedings as those 'left out' in the two-pile version. It is important to note that this *need not* be because an indeterminately placed marble indeterminately has either of the pair of effects we are tracking. Rather, the players may mutually recognize the indeterminate placement of the marble in a way that ensures the placement determinately lacks both effects. This is an important idea we will come back to. For now, I just want to flag that indeterminacies in placement do pry apart the effects we are examining without requiring a 'two-pile' set up. Moreover, it seems very unlikely that any similar game set up can avoid all such kinds of indeterminacies.

We can dispel the illusion that the added pile is really doing heavy lifting here by issuing the challenge to create four effects of a move in the context of the game which involve non-equivalent permutations of the consequences of assertion in its standard function. One can't do this *just* by adding a third pile and describing some function for it. One would also have to motivate the existence of an entirely new consequence of assertion, and show how this consequence is brought out by the use of the third pile.⁸

Also the absence of special game-driven rules governing the sorting procedure is relevant here. I'm *not* trying to make the hackneyed claim that truth-value

⁸This is not to say one can't achieve these tasks. I suspect there are indeed additional helpful ways to motivate complicating the effects of assertion through the use of my game than those I'm exploring here—for example by introducing something like 'probabilistic piles' to it. Again, in doing so one would have to pair the example with a convincing motivation for an additional function that assertion has—not just to rule possibilities in or out, but to convey likelihoods, for example.

gaps could track something like ‘misplays’. First, that analogy is inapt in important ways. It is more appropriate as an analogy for failing to assert than asserting (perhaps imperfectly) with special structured effects. Moreover the points of similarity drawn by the analogy, where appropriate, are simply too few. Not all games are like making assertions, and so calling a gap a ‘misplay’ tells us nothing on its own about what could count as a ‘misplay of assertion’ in the sense Dummett’s challenge requires.

So we should avoid the thought that interesting work is done in my analogy by artificial features of the game set up, or that the point being made is just that attempts to assert, like attempts to perform any action, can fail. If we avoid these pitfalls, the game shows how contributions to informational potential and to accuracy conditions can be pulled apart in a way that makes sense of a third status in tracking assertoric effects, which is a good start for the gap theorist. In particular, it enables her to resist Dummett’s most damaging conclusion: that the characteristic function of assertion on its own forbids the utility of a tripartite classification of assertoric effects. It also creates some space for the gap theorist to develop her positive views in a natural way. There are three tasks she has yet to perform:

- (a) The gap theorist needs to give empirical or theoretical grounds for thinking that some expressions through convention (or lack thereof) actually do have (u)-like effects.
- (b) She must explain why we should think of that behavior as defective in some way.
- (c) She must explain why a (u)-like status leads to the logical and projective behavior familiar from formal implementations of gap theories.

I’ll defer a discussion of (c) until §5 because understanding the relationship between (u)-like behavior and inference requires understanding of whether, and how, thought may be characterized by gaps. Achieving (b) is simple. After all, we characterized the effects lacked by gaps as prototypical *goals* of speakers in producing assertions. What we’ve seen is that there is an important distinction between two different ways to fall short of achieving those goals when we consider how the assertion’s operations relate to the actual world. One can conventionally rule out the actual world, or one can fail to say how things stand with it at all *while* nonetheless pronouncing how things stand with respect to other possibilities. Each involves a different kind of inaptness of assertion. So it is easy to see why gaps can sometimes be construed as a kind of defect distinct from falsehood.

Addressing (a) goes beyond the scope of this paper. Each candidate expression for gappy treatment is one whose semantics is hotly contested on both linguistic and philosophical grounds, so there is hardly hope of settling the status of any one such expression, let alone several, here. Nonetheless, it might be helpful to show why it is intuitive that most, if not all expressions to which gap theories have been applied, could be modeled with (u)-like effects.

First, consider our original problem case for views that treat gaps as grades of a single assertoric effect: indeterminacies in linguistic use. I said that if we focus on a single assertoric effect, it is often perfectly clear whether an utterance of “Chuck is bald” has that effect at a world where Chuck is a clear borderline case of baldness. When we consider a gap theory based on an amalgamation of effects, however, this poses no problems. It’s quite plausible to think that we shouldn’t aim to assert “Chuck is bald” if Chuck is a clear borderline case. It’s equally plausible to think that an assertion of “Chuck is bald” isn’t appropriately used as a conventional way to rule out the actual world either. These are *just* the conditions required to think of the assertion as gappy on the present view.

Next consider conditionals with false antecedents. Dummett correctly noted that we don’t use a conditional to rule out worlds where its antecedent is false. But it is equally wrong to *aim* at producing conditionals whose antecedents are false. Again, this is the very kind of behavior indicative of gaps on the present view. Similar remarks apply to uses of empty names. Saying “Nemo is clever” arguably isn’t used to rule out worlds where Nemo doesn’t exist. Plausibly uses of the sentence somehow presuppose that Nemo exists, and characterize how the world is—that is, rules a world in or out—as a function of how things stand with him. This doesn’t mean, however, that it is therefore a good idea to assert “Nemo is clever” if there is no such person. Again, if this is really the combination of assertoric effects the words have, they should be modeled with a truth-value gap, on my present construal, at worlds where Nemo doesn’t exist.

Similar remarks apply to the other constructions that have historically been singled out for gappy treatment. It’s worth noting that the above examples—conditionals with false antecedents and empty names—are ones which, if they are defective at all, should be characterized with *gappy* treatment rather than as failed assertions. It seems important, for example, to understand a conditional with a false antecedent spoken at worlds where the antecedent is false as ‘saying the same’ as it would at worlds where the antecedent is true. Also, even though statements using names may exhibit gaps relative to *counterfactual* circumstances where the named individual doesn’t exist, this shouldn’t mean that the assertion made at the actual world was defective in any way at all. Consequently, it might be important for semanticists to have genuine gaps alongside more standard truth-values in modeling the operation of a single assertoric act, in way that could not be captured on a theory which only allowed for ‘failures to express propositions’.⁹

⁹Though I sadly don’t have the space to treat the issue here, for those familiar with them, we now have materials to sketch a response to both ‘gap-closing’ strategies proposed in the Dummett-inspired Glanzberg (2003) to obviate gaps. The first gap-closing strategy involves pressing the distinction between compositional and assertoric content, which won’t apply to the present view as it is constructed precisely with a mind to that distinction. The second gap-closing strategy is to reconstrue assertions with gaps in a conversation’s context set as involving failed assertions. As I say here, this strategy prevents us from drawing appropriate ties of assertoric similarity between failed and successful gappy assertions in different worlds, and threatens to deprive us of natural accounts of how assertions gappy at *counterfactual* worlds in a context set might still be counted as uncontroversially true. Also, treating (say) conditionals with false antecedents simply as failed assertions may leave unexplained how

4 Gaps in Thought?

In arguing *with* Dummett I stayed agnostic as to what we might use truth-values to classify, aside from the claim that they should somehow be pertinent to assertoric effects. In the previous section, I had to make more commitments to outline the form of my positive proposal. This raises the question: did these commitments privilege some conception of truth-bearers? For example, it is often claimed that truth-values are most properly assigned to abstract entities, like propositions, which needn't figure in assertions and may act as the objects of *thought*. What is the status of gaps in characterizing thought and its objects? Can we construe gaps as applying to propositions, and not merely utterances or speech acts? Here my proposal has picked up some important commitments. Let me begin with thought.

One might think the tight connections between language and thought call for their parallel treatment, but there are key differences between the linguistic and cognitive settings which make introducing gaps into thought more difficult than it might first seem. A first thing to note is that the role of a gap, as I've given it, isn't simply to mark off a property of information. It has been to mark off a feature arising from an *act* of attempted information transfer about the actual world whose efficacy hinges on convention. To be sure, accuracy conditions can be considered as part of an abstract piece of information itself, independently of any act. But accuracy conditions are also integral to evaluating an act of assertion. Because of this, there are worries that the absence of some conventional performance will obviate the need to track accuracy assessment at a world independently of information potential.

To see this, let's consider an example where truth-values could be used to talk about information in the absence of a performance. A standard case might be in modeling the information a reasoner gathers from a piece of evidence—say a certain experience. Here the information gathered from the evidence is incompatible with certain ways the world might be. As such, when we model the information acquired we might do this with an asymmetric partition of possibilities, with one side having the now familiar ruling-out effect. And we might use what I called “accuracy inducement” to characterize the parts of the partition which don't have that effect.

Importantly, though, it's not as if the world ‘asserts’ the information modeled. Because there is nothing like an act of supplying the information, it's not clear what in the process of evidential acquisition could count as anything like a ‘sorting failure’ in my game. Who would fail to sort appropriately? Perhaps the reasoner. But what point could there be in distinguishing among the worlds the reasoner has not ruled out with the help of the evidence? No point, if we are just modeling the information acquired and no further act of supplying it. Unlike in the case of assertion, it is difficult to think of the evidence acquired as failing to be accurate information in multiple different ways.

we learn so much information about the mental state of the speaker who produces such a failed assertion—clearly something we learn on the basis of assertoric functions and not compositional roles.

Part of what I am trying to highlight with this example is that much importance in tracking contributions to accuracy conditions of what is expressed in assertion has to do with the importance of the act of asserting. It is those making assertions who are ‘on the hook’ for the information they provide, and they can be held accountable when their assertions are inappropriate, or imperfect. Accuracy conditions are an important part of what makes these assessments of an assertion possible. They ground practices of holding speakers responsible in testimony, and processes of rectification and adjustment in the event of error. It’s not clear that without a performance anything like these practices can resurface, straining the transposition of a threefold classification of effects in assertion to thought.

Now some kind of ‘performance’ might underly belief formation. After all, occurrent beliefs can be acquired at specific times, in specific circumstances, and we can be evaluated, and criticized, for our belief forming dispositions. Even if this is the case, the example of evidence acquisition is supposed to show how performance failures in belief formation and performance failures in assertion are different enough that an explanation of the purpose of a ‘third status’ in the former is liable to be very different from that in the latter. It’s not obvious that someone’s belief forming dispositions open up room to think of the worlds they rule out in acquiring their beliefs (or the worlds they don’t) in two different ways.

Earlier I said that part of what makes tracking a third status in assertion possible is that assertion involves a conventional performance. So far I have focused on the special contribution of the performative aspect of assertions. But an equally important difference between language and thought arises in considering the role of convention in speech. According to my story about the source of gaps in assertion, these gaps often arise from classificatory failures. Consider the marble game: what gives a marble a special third kind of status is that it can’t easily be classified as having either of the two more standard effects in an ‘assertion’. An aspect of any assertoric practice which is essential for this kind of classification failure is the conventional association of informational effects with a performance—say uttering a particular string of sounds, or making a gesture. Informational effects of speech are *encoded* in an act, and *decoded* from it by interpreters.

These features of information transfer are important because classification failures are often *made possible* by the encoding and decoding processes. It is primarily because there is a non-trivial task of coordination on interpretation that there is room for mistakes or uncertainty in what kind of status a particular string of sounds, say, is ‘supposed’ to have.

The encoding and decoding processes necessarily present in assertion do not obviously have any analogs in thought. Even if there are processes in cognition whereby beliefs are stored in our heads in some recognizably linguistic form—an idea that is already highly controversial—it’s not obvious that there is a special problem of interpretation that we face when retrieving and manipulating the sentences in this language of thought. In fact there is every reason to think that processes of interpretation must stop somewhere, and that the stopping

point would occur at or before the manipulation of this hypothetical mental symbolism.

So to sum up: my story about the source of gaps in assertion relied on the fact that assertions were conventional performances. That speakers, as performers, take responsibility for the information they supply contributes to the importance of tracking accuracy conditions independently of information potential. Moreover, it is largely because such acts are conventional, and involve processes of encoding and decoding information, that room for failure to appropriately track the assertoric effects of an assertion could arise. Without room for the relevant kind of failure in thought, there is no longer any clear role for the use of gaps.

I haven't argued that no sense at all can be made of the utility of gaps, or other kinds of third statuses, in describing cognition and cognitive acts. But I have argued that *if* we are to make sense of such gaps, our account of their existence and utility will depart from the one I've supplied so far.

I regard this as a very welcome conclusion. The resulting view is one on which assertions that generate gaps *need not* be indicative of defects in the mental states of the speakers who produce them. This is arguably *needed* to deal with many allegedly gappy utterances in a satisfactory way. Take conditionals with false antecedents. Suppose that p is false, though if p were the case, then q would be as well. Tamara hasn't settled whether or not p , and so utters the indicative conditional "if p then q ". When we've settled that p is false, we may not be able to comfortably classify Tamara's utterance as straightforwardly true, or straightforwardly false. But was there something 'defective' about Tamara's belief state, or even something false that she believed? Not obviously. It seems that everything I've said so far is compatible with Tamara only having *imprecise* beliefs. Perhaps Tamara merely hasn't settled yet whether p is true, and were she to do so she would only *acquire* beliefs, all of which are true. To take another example, speakers who use non-referring names needn't have anything like a 'gap' in their mental state. They might simply have several clearly false beliefs—about what sorts of things do and don't exist, and about what their words refer to. It's only their *assertion* which exhibits a special kind of defective character.

So far from being a disappointing result, it should be viewed as a welcome flexibility in the present view that it doesn't immediately entail that gaps in assertion must result from defective mental states. We can still allow for the existence of such states if we like, just in different ways than I have done for the defects in assertion.¹⁰

So much for belief. What of propositions? This issue here is much more complex than for the structure of mentality, since whether the present theory allows gaps as features of propositions depends on our conceptions of their nature and theoretical role. For example, if propositions are identified with certain bodies of information, then as I've argued there is probably no room

¹⁰A natural way to introduce trivalence into the characterization of mentality might be, for example, to model forms of agnosticism and indifference—states which in turn seem to have no natural analog in the operation of assertion.

to make sense of gappy propositions. If propositions are conceived instead as theoretical tools used to model regularities in assertoric functions, propositions may be trivalent. This is one among several ways of cashing out the claim that propositions act as ‘objects of assertion’. Such propositions might even be of use in characterizing thought indirectly: speakers may use trivalent entities to characterize a belief state in a belief report by effectively sorting among belief states which leave open only possibilities at which the proposition exhibits bivalence. This is, in turn, one among several ways of thinking how propositions act as ‘objects of thought’.¹¹ But if, by contrast, propositions are considered to be *part* of a mental state, and not merely an external means of characterizing a belief state, trivalence of the kind I’ve defended will again be more difficult to attribute to propositions.¹²

This is to say that my defense of gaps *may* have chalked up some significant commitments for how we should construe propositions in ways that it definitely has for how we should construe the relationship between gaps and mentality. Are these theoretical commitments detrimental in any way? It’s not, or not obviously, so I think, if regardless of how we construe the use of trivalence in characterizing assertions, that trivalent scheme will be able to play the roles philosophers have traditionally demanded of it in the linguistic setting. I’ll now argue this is the case.

5 Projection and Inference

Many have supposed that we could argue for the existence of gaps by pinpointing a special kind of projection behavior, and that the gaps uncovered in this way would transform our conceptions of logical consequence. Dummett ably challenged both of these ideas. He showed why we should keep compositional roles separate from assertoric effects without special arguments connecting them. He also gave powerful reasons to doubt the utility of any tripartite characterization of the content we make inferences to and from.

My characterization of gaps, however, gives hope of turning the tables back in favor of gap theorists. Getting a firm grip on the nature of gaps actually *strengthens* the connections between defects in assertoric function, strong projection behavior, and the disruption of inferential schemes.

Consider projection first. On my story, gaps can arise due to certain kinds of classificatory failures. An assertion will have an effect-at-a-world characterized by a gap when there is sufficient uncertainty or reluctance among language

¹¹For an elaboration of the kind of relationship between propositions and belief states I’m drawing on here see Stalnaker (1987), and especially the discussion of ‘derived context sets’ as tools to test when a proposition can felicitously be used to characterize a belief state.

¹²This discussion bears on the question of whether my third value is a ‘genuine truth-value’ if it can’t be used to characterize propositions. I’m happy to concede that my third value might not be a ‘genuine truth value’ in this sense, provided it is acknowledged that the tripartite classification of assertoric effects I’ve highlighted is useful in precisely the ways that philosophers have assumed that a tripartite allotment of ‘genuine truth values’ are. I argue this in the next section.

users in classifying the assertion's effects at the relevant world with the effects given by (t) or (f). It seems plausible to suppose that the effects of assertions involving compounds of sentences are determined by operations on the effects of their constituents. That is, the truth-value assignment we use to classify assertoric effects *typically* also plays the role of classifying compositional effects—at least for simple logical connectives like “and” or “not”. If this is right, then there is every reason to think that the circumstances leading to classificatory failures at the level of assertoric effects translate into classificatory failures of compositional role. That is, parts which would exhibit gaps *tend* to produce gaps when embedded, which is just the compositional behavior we want gaps to have.

This might not be terribly surprising. But, importantly, this source of strong projective behavior does *not* show that it is unavoidable—only that we should expect it to be the default behavior. There are several ways in which projection would be blocked. This could occur, for example, when speakers coordinate on the existence of the classification failures generating gappy behavior. Consider an analogy: Two umpires witnessing an unlikely event during play in baseball might each be at a loss as to what the ruling should be on that play—on what the rules of baseball prescribe for this circumstance. They might nonetheless each *know* of the other that they too are at a loss. In this case there is a kind of classificatory failure among the two umpires, and perhaps others, as to how a particular play should be ruled. But there might be near complete coordination on whether the case is one where there is, or will be, such classificatory failure.

There seems to be no bar to similar cases arising for assertions. In fact, in giving this paper I have been *relying* on the fact that we can sometimes agree on when certain expressions exhibit the defects characteristic of gaps. When such agreement arises, there seems to be no bar to our recognizing the classificatory failure and its potential effects. There is equally no bar to creating concepts to describe the existence of these classificatory failures, and their effects, whether it be in baseball or in assertion. So on the story I've supplied we have this immediate result about the projective behavior of gaps:

Truth-value gaps generally project under embeddings unless:

- (i) there is suitable coordination among language users on the behavior of the embedded expression responsible for the gap, and
- (ii) there are compositional devices whose explicit function is to somehow exploit the presence of such gaps for communicative purposes (e.g., by reporting their existence).

One might doubt that the simplest kinds of logical connectives are tailored to the kind of special task in (ii). If so, gap theorists are vindicated in thinking that strong projective behavior counts as evidence for philosophically interesting truth-value gaps.

Let's turn now from projection to logical consequence. A logical consequence relation typically tracks how syntactic structure mediates legitimate inference. But what *is* an inference or a legitimate inference in the first place? A prevalent

view takes inference to be a kind of mental state transition, for example from one belief state to another. A belief state transition counts as legitimate if the transition ensures the second belief state will be true provided the first one is.

Now, if good inference is a property of belief state transitions, why do we often speak of sentences as standing in inferential relations? Naturally because sentences can be used to express the relevant belief states, or characterize them. So it can make sense to say that a sentence transition involves good inference in a kind of derivative way: it is effectively to say that the corresponding transition between the belief states that the sentences express, or would characterize, is a transition involving good inference.

Here is where a caveat must be added if truth-value gaps arise in the ways I have claimed, because gaps complicate the relation between sentences expressing or characterizing beliefs and the beliefs themselves. As I stressed in §4, though gaps arise in characterizing assertions, it is not obvious that they can be transposed in any direct way to the description of thought. Thus it is not yet clear what one could mean in stating that a transition between sentences with gappy behavior is a good one. It will not do simply to say that it means that the transition between the belief states expressed or characterized by the sentences is a truth-preserving one. After all, the question of exactly what belief state a gappy proposition is used to represent, or express, if any, is so far up for grabs.

It is tricky to do more here than point out this initial difficulty. Part of the problem is that we are encroaching on complex issues connected with the semantics of belief ascriptions. Can we ever use an entity marked by gappy behavior in assertion to characterize a belief state which does not exhibit such gaps (or gaps of the same kind)? If so when, and how? This is a topic I cannot hope to treat in any detail here. What we can say immediately is that accommodating gaps is *bound* to transform our conceptions of the language-bound representation of logical inference, if only in how we state, and understand, the relationship between the standard sentence-to-sentence transitions that are the stock and trade of formal logicians, and the belief-to-belief transitions which give our logical apparatus its point.

For now, though, we can sidestep some of these questions about belief ascriptions by moving to a much simpler conception of inference worth characterizing in abstraction from questions about thought. This is a notion of inference restricted to sentence transitions: such sentence transitions are good ones when truth is preserved in virtue of their logical form. If we abstract from sentences' customary role in characterizing appropriate transitions in thought and focus on, say, assertability, then we would expect gaps to play a role like falsehood in regulating logical inference simply by stipulation. Transitions from truthful to gappy sentences are clearly bad ones, and transitions from gappy sentences to any other kind are good, if ever, only vacuously. And though it is clear that gaps will have a role just like falsehood in regulating *these* kinds of inferential relations, we already have seen good reasons to think that gaps have a much stronger projection behavior than falsehood. Because of this we should expect the presence of gaps to transform the inferential behavior of standard logical connectives. In particular, truth-preserving inferential logical form will

be harder to come by which means, as gap theorists have long claimed, that logical form will be a much less useful means for tracking good inferences than is often supposed.

This is all just to say that the story about the source of truth value gaps I have supplied vindicates the two standard roles that gaps have had to play in philosophy, compositional semantics, and logic. Gap theorists have rightly been criticized for leaving the basis for these two roles, and their relation, undeveloped. But showing how a natural account of assertion gives rise to a phenomenon with both features gap theorists sought strengthens the sense that they were tracking instantiations of the same phenomenon, even if with a loose or imperfect conception of how it arose. It also gives hope that the utility of these two roles of truth-value gaps in application to many philosophical problems is also safeguarded, though I won't be able to explore any of these applications here.

References

- E. Camp (2004). 'The Generality Constraint and Categorical Restrictions'. *The Philosophical Quarterly* **54**(215):209–231.
- M. Dummett (1978). 'Truth'. In *Truth and Other Enigmas*, pp. 1–24. Harvard University Press, Cambridge, MA.
- M. Glanzberg (2003). 'Against Truth Value Gaps'. In *Liars and Heaps: New Essays on Paradox*, pp. 151–194. Clarendon Press, Oxford.
- G. Priest (2006). *In contradiction: a study of the transconsistent*. Oxford University Press, USA.
- R. Stalnaker (1978). 'Assertion'. In *Context and Content*, pp. 78–95. Oxford University Press, Oxford.
- R. Stalnaker (1987). 'Semantics for Belief'. *Philosophical Topics* **15**(1):177–190.
- R. Stalnaker (2002). 'Common Ground'. *Linguistics and Philosophy* **25**(5):701–721.
- P. F. Strawson (1950). 'On Referring'. *Mind* **59**(235).