

10

Embedding Denial

David Ripley

10.1 Introduction

10.1.1 A puzzle about disagreement

Suppose Alice asserts p , and the Caterpillar wants to disagree. If the Caterpillar accepts classical logic, he has an easy way to indicate this disagreement: he can simply assert $\neg p$. Sometimes, though, things are not so easy. For example, suppose the Cheshire Cat is a paracompletist who thinks that $p \vee \neg p$ fails (in familiar (if possibly misleading) language, the Cheshire Cat thinks p is a *gap*). Then he surely disagrees with Alice's assertion of p , but should himself be unwilling to assert $\neg p$. So he cannot simply use the classical solution. Dually, suppose the Mad Hatter is a dialetheist who thinks that $p \wedge \neg p$ holds (that is, he thinks p is a *glut*).¹ Then he may assert $\neg p$, but it should not be taken to indicate that he disagrees with Alice; he doesn't. So he too can't use the classical solution.

The Cheshire Cat and the Mad Hatter, then, have a common problem, and philosophers with opinions like theirs have adopted a common solution to this problem: appeal to denial.² Denial, these philosophers suppose, is a speech act like assertion, but it is not to be understood as in any way reducing to assertion. Importantly, denial is something different from the assertion of a negation; this

¹ There's a pesky feature in terminology here. A 'paracompletist' is typically taken to be someone who thinks that some instances of excluded middle fail (see e.g. Field 2008), but a 'paraconsistentist' need not think that some contradictions are true. That's a 'dialetheist'. So parallel terms don't pick out parallel points in philosophical space (and the position dual to paraconsistentism doesn't even have a standard name). Nevertheless, here I follow standard, if suboptimal, terminology.

² For details on specific ways of doing this, see for example Priest (2006a), Field (2008), Beall (2009), and Parsons (1984). The description I offer here of denial is common to all these authors. Note that some of these authors prefer to speak of *rejection* (an attitude) rather than *denial* (a speech act); I don't think the difference matters for my purposes here (simply swap in 'acceptance' where I say 'assertion'), and I will henceforth ignore it.

is what allows it to work even in cases where assertion of negation does not. Just as importantly, denial must express disagreement, since this is the job it's being enlisted to do.

Let's consider in this light the case of the Cheshire Cat. He disagrees with Alice's assertion of p , but is unwilling to assert $\neg p$. He can indicate his disagreement simply by denying p . A paracompletist who denies p is not committed to asserting $\neg p$; in fact, the Cheshire Cat would deny $\neg p$ as well. Dually, a dialetheist who asserts $\neg p$ is not committed to denying p ; the Mad Hatter can assert $\neg p$, and we need not take that as indicating his disagreement with Alice. If he were to go on to *deny* p , however, we'd understand him as disagreeing.

Nothing about this story, let me hasten to note, undermines the Caterpillar's strategy for expressing disagreement. Since the Caterpillar accepts classical logic, we can take his assertion of $\neg p$ as committing him to a denial of p . Thus, his assertion of $\neg p$ is still successful in expressing disagreement with Alice, so long as he is committed to classical logic. This theory about denial, then, has the virtue of allowing us to make sense of a range of disagreements, even when those disagreements occur between characters with different logical commitments.³

10.1.2 *The denier paradox*

In order to work as I've described it, denial must satisfy both of the following principles:

D-exclusivity: Asserting p and denying p are *incompatible* speech acts; they rule each other out (otherwise we could have denial without disagreement)

D-exhaustivity: Having a settled opinion about p requires being either willing to assert p or willing to deny p (otherwise we could have disagreement without denial)⁴

These are parallel to the classical principles of \neg -exclusivity (that p and $\neg p$ can't both be true) and \neg -exhaustivity (that at least one of p , $\neg p$ must be true). Since both the paracompletist and the dialetheist give up the conjunction of these negation principles to avoid paradox, one might naturally worry that invoking a notion like denial, one that satisfies the analogues of the negation principles, might lead right back to paradox.

³ For more on this particular virtue, see Restall (2013).

⁴ In ordinary life, one might have a settled opinion about p without being willing to say much at all about it—maybe p is a secret, or too long or boring to bother with. Here, when I talk about willingness to assert or deny, I mean *in a complete statement of one's theory*. (This is one reason why the difference between assertion/denial and acceptance/rejection can be ignored here.)

Here's the threat: consider a sentence (or a content) δ (we'll call it the *denier*) such that asserting δ is equivalent to denying δ . That is, asserting it commits us to denying it, and denying it commits us to asserting it. If we assert it, we're committed to denying it, but that's unacceptable (by D-exclusivity). If we deny it, we're committed to asserting it, but that's unacceptable (by D-exclusivity). So we'd better neither assert it nor deny it. But since this is (must be!) our settled opinion, this is unacceptable too (by D-exhaustivity). So we're in trouble no matter what we do.

This is only a threat, however, and not genuine trouble, so long as there is no such δ . Paracompletists and dialetheists who appeal to denial must take care that their theories do not countenance such a δ , and indeed they do take such care.

Avoiding such a δ requires avoiding certain operations on content. Suppose we were to consider a unary content operator D such that to assert DA , for any content A , is equivalent to denying A . Then there would be a δ as specified above: simply let $\delta = DT\langle\delta\rangle$, where $T\langle\rangle$ is an intersubstitutable truth predicate, augmented with the usual sort of self-reference-allowing naming device.⁵ So it's crucial to the tenability of our denial story that there be no such operator D .

In section 10.2, however, I'll argue that we need an operator like D to make full sense of denial. Once we have D (and therefore δ), the denier paradox pushes us to give up either D-exclusivity or D-exhaustivity; but then the initial puzzle about disagreement cannot be solved in the way I've outlined. I consider other possible ways to solve the puzzle about disagreement in section 10.3 and section 10.4.

10.1.3 Does denial exist already?

In Parsons (1984), denial is assimilated to 'metalinguistic negation' in the sense of Horn (1989). Metalinguistic negation is the sort of negation occurring in, for example, 'Bryce doesn't have THREE helicopters; he has FOUR'. It seems distinct from an ordinary negation because the utterer of such a sentence doesn't mean to say that Bryce doesn't have three helicopters. In fact, for the sentence to be true, Bryce *has* to have three helicopters. He just also needs to have a fourth one. The utterer of such a sentence is rejecting the *appropriateness of an assertion of* 'Bryce has three helicopters', not because the sentence is false, but instead because it would be misleading.

⁵ Both Field (2008) and Beall (2009) advocate for such a truth predicate in their respective frameworks. While Priest (2006a) and Priest (2006b) argue against truth's being intersubstitutable, there's no reason not to define an intersubstitutable predicate in Priest's framework. He might insist it's not truth, but for its role in constructing δ it doesn't matter whether or not it's truth; it only matters that the predicate be intersubstitutable.

Parsons focuses on this use of metalinguistic negation—to object to a sentence without necessarily committing oneself to an assertion of the sentence’s negation—and says that this is our real-world denial. According to Parsons, we use the word ‘not’ in two different ways (Horn (1985) calls this a ‘pragmatic ambiguity’): sometimes it modifies the content of our sentence, and sometimes it leaves the content alone, but indicates that the speaker rejects the sentence on some grounds or other. (In the above example, the grounds are grounds of misleadingness.) When that rejection is based on the sentence’s content, we have a denial, according to Parsons’s theory.

There’s a serious problem with understanding denial in terms of metalinguistic negation, however. Denial is supposed to be a certain type of speech act, and speech acts are something we do with a content. We can’t build contents with speech acts as parts. However, as Geurts (1998) points out, metalinguistic negation has no difficulty embedding into larger contents; the following sentences are perfectly fine:

- (1) If Bryce doesn’t have THREE helicopters but FOUR, then he has one more helicopter than I thought.
- (2) Mary thinks Bryce doesn’t have THREE helicopters but FOUR.

Geurts points this out as part of an attack on Horn’s ‘pragmatic ambiguity’ thesis, but it tells just as strongly against Parsons’s analysis of denial as metalinguistic negation.

In fact, I don’t know of any phenomenon studied outside the realm of philosophical logic that could fill the theoretical role occupied by denial in our philosopher’s theories.⁶ I’m not sure, though, that that in itself is a problem for these theories. After all, denial in the present sense only needs to be invoked (that is, it only differs from asserting a negation) in the presence of gaps or gluts. If most speakers, most of the time, are not worried about gaps or gluts (as seems plausible), then we shouldn’t necessarily expect there to be an obvious ordinary-language correlate of denial. Those of us who are worried about gaps and gluts may simply have to *introduce* a new sort of speech act to make plain our meanings, and I don’t see that there’s anything to stop us from doing this.

Indeed, Field (2008, 96) seems to treat denial roughly along these lines, proposing that saying something ‘while holding one’s nose’ or writing something in a certain ugly font could serve, if we so chose, as marks of denial. There’s obviously

⁶ Geurts (1998) mentions another phenomenon traveling under the name of ‘denial’, but it is clearly unrelated. I’ll mention it again in n. 12.

something arbitrary about these choices, but arbitrariness should be no barrier here.

Of course, if we are *introducing* denial rather than *discovering* it, it's important to be clear about just what we're introducing. Which features does this speech act have? Which does it lack? I think when we try to get clear about these questions, we will see that a speech act that behaves like denial does ought to have a corresponding operation (like D) on content. And, as we've seen, that leads to revenge paradox. Section 10.2 spells this out.

10.2 Denial in the Image of Assertion

Denial, for these theorists, is not any kind of assertion, but it's at least something *like* assertion.⁷ In fact, it's more like assertion than either of them is like questioning or ordering. Both assertion and denial are *informative*; they attempt to tell us something about the way things are. This suggests that, to learn about denial, we should look at theories about assertion, and make the appropriate modifications.

In this section, I'll do just that. First, I'll look at norms assertion is sometimes alleged to fall under, and point out some trouble for stating the appropriate parallel norms on denial. Second, I'll look at uses of truth together with assertion to express agreement, and point out some trouble for using truth to express agreement in the presence of denial. Third, I'll point to a certain sort of question about priority that is a sensible question to ask about assertion, and point out some trouble for stating a parallel question about denial. All three of these troubles can be solved if we have a D operator in the language. (Of course, as we've already seen, D brings its own troubles. Nothing in this section will alleviate any of those. The purpose of this section is rather to make clear the troubles that arise from a *lack* of D.) Fourth, I'll show how to generalize a Stalnakerian theory of assertion to encompass denial, and point out that this allows us simply to define D. All these considerations, taken together, provide a compelling argument for having D around. I close this section with some reason to think having D around isn't totally hopeless, despite the risk posed by the denier paradox.

10.2.1 Norms governing denial

Assertion is often taken to be subject to certain norms. Just what these norms are is a matter of some dispute; here I'll consider two options:

⁷ For example, Field (2008, 74) offers: '[W]e should regard acceptance and rejection as dual notions. And how exactly one thinks of rejection will depend on how one thinks of the dual notion of acceptance.'

(Assert-T) Assert A only if A is true

(Assert-K) Assert A only if you know A

10.2.1.1 (ASSERT-T)

Let's start with (Assert-T). Suppose this is a norm that governs assertion. The natural question is: what is the corresponding norm governing denial? Here are two possibilities:

(Deny-F) Deny A only if A is false

(Deny-NT) Deny A only if A is not true

Neither of these, though, should be acceptable to the proponent of denial. Let's consider (Deny-F) first. If (as is standard) we take falsity to be truth of negation, then (Deny-F) tells us to deny something only when its negation is true. But the Cheshire Cat couldn't then use denial to indicate his gappy take on p , not without (by his own lights, anyway) flouting (Deny-F). On the other side of the coin, the Mad Hatter ought to find (Deny-F) unacceptably weak; he thinks there are plenty of false things (Alice's p among them) that ought not be denied. What's more, it's not that they shouldn't be denied for reasons of politeness or some such; it's that one would be *mistaken* to deny them, in the same way one is mistaken in asserting something that isn't true. Such cases ought to be covered by a norm governing denial parallel to (Assert-T), but (Deny-F) fails to do the work.

If we turn to (Deny-NT), things are little better. In fact, if we have an intersubstitutable truth predicate, things are no better at all, since there is no difference between A 's being false (that is, $T(\neg A)$) and A 's not being true (that is, $\neg T(A)$), and so no difference between (Deny-F) and (Deny-NT). This already rules out (Deny-NT) for Field-style paracompletists and Beall-style dialetheists.

As above, however, not every theorist who appeals to denial accepts that truth is intersubstitutable. In particular, Priest-style dialetheists and dual sorts of paracompletist do not. (There are probably no actual instances of Priest-duals, but there certainly could be.) For example, Priest takes the simple liar— $\lambda' = T(\neg \lambda')$ —to be false (he asserts $T(\neg \lambda')$), but he doesn't take it to be untrue (he denies $\neg T(\lambda')$). His dual would do just the reverse: assert $\neg T(\lambda')$, but deny $T(\neg \lambda')$. Indeed, when untruth and falsity come apart in either of these ways, (Deny-NT) outperforms (Deny-F): Priest won't deny λ' , and his dual will. This is in accord with (Deny-NT), but not (Deny-F).

Unfortunately, this works as nicely as it does only because the simple liar is a special case. When we turn to the *strengthened* liar— $\lambda = \neg T(\lambda)$ —things don't work as cleanly. This is because Priest and his dual agree that falsity and untruth

can't come apart for λ . Suppose λ is false. Then we can argue that it is untrue as follows:

1. $T\langle\neg\lambda\rangle$ Assumption
2. $\neg\lambda$ 1., Release⁸
3. $\neg\neg T\langle\lambda\rangle$ 2., Substitution
4. $T\langle\lambda\rangle$ 3., Double Negation
5. λ 4., Release
6. $\neg T\langle\lambda\rangle$ 5., Substitution

On the other hand, suppose λ is untrue. Then we can argue that it is false by running the above argument backwards (each step is valid in the other direction as well). Thus, when it comes to λ , we lose the distinction between (Deny-F) and (Deny-NT). This is to (Deny-NT)'s detriment; Priest and his dual should have the same complaints about (Deny-NT) applied to λ as the Mad Hatter and Cheshire Cat had about (Deny-F) in general: it is unacceptably weak.

10.2.1.2 PRIEST AND (DENY-NT)

This is enough to indicate the problems that the paracompletist and dialetheist should have with (Deny-F) and (Deny-NT). Neither can work as a parallel to (Assert-T). However, there is a potential objection to this conclusion, suggested by remarks in (Priest 2006a, section 6.5). (I owe this observation to Priest (pc).)

Priest considers endorsing something like (Deny-NT) as a principle governing rejection. His principle: 'One ought to reject something if there is good evidence for its untruth'. Ignoring the difference between rejection and denial, this is an evidentially flavored version of (Deny-NT). Rather than consider this principle directly, I'll consider how Priest's defense of it would apply to (Deny-NT). (This affects nothing of substance.)

Recall that according to Priest the strengthened liar λ is both true and untrue. As such, (Deny-NT) comes into conflict with (Assert-T) in the case of λ ; according to these principles we should both assert it and deny it. Priest's suggestion is that this may simply be a rational dilemma; that in virtue of (Deny-NT) and (Assert-T) both holding, we are under obligations that we cannot fulfill.

From one point of view, this is very close to the approach I'll recommend in section 10.4.2. However, it is more radical than Priest acknowledges. For example, consider a discussion between Priest and the (paracompletist) Cheshire Cat about

⁸ I assume that $T\langle\rangle$ validates two inferences: Capture (from A to $T\langle A\rangle$) and Release (from $T\langle A\rangle$ to A). This is weaker than the assumption of intersubstitutivity, and is acceptable to Priest and his dual.

the strengthened liar. The Cheshire Cat denies it, while Priest asserts it. But Priest can have no leverage with which to criticize the Cheshire Cat; by (Deny-NT) the Cheshire Cat is behaving as it ought to. Of course, the Cheshire Cat is also behaving as it oughtn't (by (Assert-T)); but so too is Priest, by (Deny-NT). On the rational-dilemma approach, Priest has no firmer grounds for criticism of the Cat than he has for self-criticism.

Thus, (Deny-F) and (Deny-NT) should leave the orthodox dialetheist and paracompletist alike unsatisfied as parallels to (Assert-T). But (Assert-T) is not the only norm of assertion on offer. Perhaps we can do better finding a parallel to (Assert-K)?

10.2.1.3 (ASSERT-K)

Here are two that won't work:

(Deny-KF) Deny A only if you know A is false

(Deny-KNT) Deny A only if you know A is not true

These fail for the very same reasons as (Deny-F) and (Deny-NT), respectively.

There is another strategy we have available here, though. Parallel to the distinction between assertion and denial (speech acts), dialetheists and paracompletists often draw a distinction between acceptance and rejection (attitudes). By asserting, one indicates acceptance, and by denying, one indicates rejection. Rejection, they say, cannot be understood as acceptance of negation, for the same reasons that denial can't be understood as assertion of negation.

This suggests a more general strategy of bifurcation. To construct a norm on denial parallel to (Assert-K), maybe we shouldn't try to build up anything workable out of knowledge and negation. Instead, maybe we should postulate a new attitude, *knowledge_D*, parallel to knowledge in the same way rejection is parallel to acceptance, and denial to assertion. You might worry that this bifurcation strategy is a bit clumsy. But it turns out to face more trouble than that.

If we understand knowledge as justified true belief, we should probably understand *knowledge_D* as justified D-true rejection (where A is D-true iff $T\langle DA \rangle$). This, of course, requires us to have the paradoxical D operator in our language, and so is not acceptable. Fortunately, there's no need here to accept the JTB account of knowledge; we can simply take both knowledge and *knowledge_D* as primitives. Then the proposed norm on denial is:

(Deny-K_D) Deny A only if you know_D A

Unfortunately, *knowledge_D*, even without the D operator on its own, is problematically paradoxical. Consider the know_Der paradox— $\kappa_D = \langle \kappa_D \rangle$ is know_Dn.

Suppose someone know_D s κ_D . Then κ_D is true; but true things can't be know_D n.⁹ We can't assert our supposition, so, by D-exhaustivity we can deny it; in fact, we now know_D our supposition. But our supposition just was κ_D , so we know_D κ_D . We must both assert and deny our supposition, and this is no good, by D-exclusivity.

So while (Deny- κ_D) forms a fine parallel to (Assert-K), it mires us in paradox. It does not seem that either (Assert-T) or (Assert-K) has a parallel that can apply to denial without trouble: the proposed parallels either fail to be properly parallel, are unstatable without D, or mire us in D-style paradox on their own.

10.2.1.4 FINDING THE RIGHT NORMS

Now, suppose we have our denial-embedding content operation D around. Then we can formulate two new principles:

- (Deny-D) Deny A only if DA is true
- (Deny-KD) Deny A only if you know DA

These are not only parallel to (Assert-T) and (Assert-K); they are instances! Remember, denying A is equivalent to asserting DA. So (Deny-D) amounts to: Assert DA only if DA is true, an instance of (Assert-T). Similarly, (Deny-KD) amounts to: Assert DA only if you know DA, an instance of (Assert-K). So the presence of D allows us to state norms on denial in a simple and straightforward way. These norms are, as we want, parallel to our norms on assertion.

Of course, just as with knowledge_D , D causes problematic paradox. But it allows us to give a straightforward theory of norms governing denial, which we do not seem to be able to do without paradox. What's more, it allows us to avoid the clumsy bifurcation strategy in general. knowledge_D is just ordinary knowledge of a D-content; rejection is just acceptance of a D-content; etc.

10.2.2 Agreement and generalizations

One use we have for the truth predicate is to make certain generalizations:

- (3) Everything Alice says is true
- (4) If everything Alice says is true, I'll eat my hat

The motivations have been stressed by Field (2008) and Beall (2009), who use them to argue for an intersubstitutable truth predicate. However, although their

⁹ On a dialetheist line, one can know the negation of a true thing (if that negation is also true), but knowledge_D and truth remain incompatible—that's the point of distinguishing knowledge_D from knowledge of negation.

theories provide fully intersubstitutable truth predicates, those predicates don't quite work to generalize in the way they seem to hope, because of denial.

If Alice only asserts, and never denies, then (3) and (4) serve their purpose well. But suppose that Alice has asserted some things and denied some others, and that Humpty Dumpty is in full agreement with Alice. That is, Humpty Dumpty is willing to assert everything Alice asserted, and willing to deny everything Alice denied. How can he indicate this?¹⁰ He might try something like (5):

(5) Everything Alice said is true

There are two readings we can give to (5), depending on how we interpret 'says'. Unfortunately, neither reading gives us what we want; on neither reading does (5) express agreement with Alice.

Suppose that Alice said something iff she either asserted it or denied it. Then (5) clearly does not convey agreement; if Humpty Dumpty thinks that Alice denied some true things, then he *disagrees* with her. So suppose Alice said something iff she asserted it; she does not count as having said the things she denied. Then (5) still doesn't convey agreement; Humpty Dumpty might agree with all of Alice's assertions, but think that some of her denials were mistaken, and still be willing to assert (5).

We might try some simple modifications:

(6) Everything Alice asserted is true, and everything she denied is false

(7) Everything Alice asserted is true, and everything she denied is not true

But these modifications get us nowhere. For the same reasons that 'false' and 'not true' would not serve to solve the initial disagreement problem—remember, only denial would do—they cannot solve the agreement problem.

Suppose Alice denies q . If the Cheshire Cat is a paracompletist who thinks that $q \vee \neg q$ fails, then he agrees with her. But he will not assert either (6) or (7). If the Mad Hatter is a dialetheist who thinks that $q \wedge \neg q$ holds, then he disagrees with her. But he will still be willing to assert both (6) and (7).¹¹ Thus, there seems to be no way to express agreement with those who deny things as well as asserting.

Now, suppose we have D in the language, and consider (8):

(8) Everything Alice asserted is true, and everything she denied is D -true

¹⁰ He can say 'I agree totally with Alice', but this will not behave properly in embedded contexts; try it in (4) to see the trouble.

¹¹ As we saw in section 10.2.1, the difference between falsity and not-truth depends on a non-intersubstitutable truth predicate. If there is such a difference, there might be some cases where (7) could work; but it still cannot work in full generality. For example, if q in the above example is the strengthened liar, it cannot work.

(For A to be D -true, recall, is for DA to be true; D -truth is thus to D as falsity is to \neg .) If truth is intersubstitutable, this expresses agreement with Alice; to assert (8) is to commit to asserting everything Alice asserted and denying everything she denied. (If truth is not intersubstitutable, this might not quite be the case, but the desire to use truth to express agreement motivates an intersubstitutable conception; it shouldn't surprise us that that remains the case here. And, as above, we can add an intersubstitutable predicate for these purposes, even if we don't think that predicate is a *truth* predicate.)

10.2.3 Questions of priority

(9) and (10) are truisms about conjunction:

- (9) If you're committed to asserting both conjuncts, you're committed to asserting their conjunction
- (10) If both conjuncts are true, their conjunction is true

One might explain (9) in terms of (10), explain (10) in terms of (9), or take both to stand on their own. Whatever route you're tempted by, there's clearly at least a question to be answered here. Are constraints on assertion like (9) prior to, posterior to, or unrelated to, constraints on truth like (10)?

Since this makes sense as a question about assertion, it ought to make sense as a question about denial. There's trouble, though. (9) has a parallel involving denial:

- (11) If you're committed to denying at least one conjunct, you're committed to denying the conjunction

Unfortunately, there is no corresponding parallel to (10). (Again, we can try to build such a parallel using falsity or not-truth, and again, such efforts will fail, for the same reasons as above.) So what is a perfectly sensible question about assertion cannot be stated as a question about denial.

That is, of course, unless we have D around. Then, parallel to (10), we get (12):

- (12) If one conjunct is D -true, the conjunction is D -true

We can sensibly ask the same questions about the relation between (11) and (12) as we asked about (9) and (10). Once again, it is only in the presence of D that assertion and denial are really parallel.

We've thus seen that three theories of assertion—to do with norms, interaction with truth and agreement, and questions of priority—*have no parallel* relating to denial, unless D is in our language. For all I've said here, we might give a theory of denial in a different way, a way that doesn't hold it to parallel our theories of assertion. But that is not how denial is understood in Priest (2006a), Field (2008),

or Beall (2009). These authors take denial to be parallel to assertion; and this simply cannot be so unless D is around.

10.2.4 Stalnakerian denial

There's a quite natural way to understand denial on a Stalnakerian picture of pragmatics. But once denial is understood that way, we can simply *define* D. Here, I explain.

In the framework of Stalnaker (1978), each stage of a conversation is associated with a set of possible worlds—the *context set* for that stage. Roughly, we can understand the context set as the set of possible worlds that are still live possibilities for the conversation. (Note that even if some (or even all) of the participants in a conversation don't personally consider a possibility live, it can still be live for the conversation, and vice versa.) The context set at any given stage constrains which conversational moves are acceptable, and how they will be interpreted. In turn, conversational moves change the context set in certain predictable ways.

In particular, assertion works by ruling out certain possibilities from the context set. When someone asserts A, the context set shrinks; any world at which A fails to hold is removed.¹² If there were no worlds at which A fails to hold left in the context set, then the assertion is infelicitous. This may trigger reinterpretation or censure of various sorts.

This can straightforwardly be extended to denial: when someone denies A (and the denial goes unchallenged), the context set shrinks; any world at which A holds is removed. This not only gives us a clear picture of denial as parallel to assertion, but it also allows us to understand what it is that assertion and denial have in common as opposed to, say, questioning. Both assertion and denial shrink the context set; this is how they are *informative*. They help us narrow down possibilities for how things might be.

The Stalnakerian picture is attractive on its own, and its extension to denial has some nice features. But it leads directly to embeddable denial. Call a move in a conversation *appropriate* iff it does not rule the world in which it is made out of the context set. Then we can build an embeddable denial: just let DA be: 'It would now be appropriate to deny A'. Asserting DA would thus rule out a world *w* iff a denial of A fails to be appropriate in *w*—that is, iff a denial of A would rule out *w*. Thus, asserting DA amounts to denying A; this was our condition on embedding denial.

¹² At least if the assertion goes unchallenged. We can ignore this caveat here, but it's worth noting that some authors (e.g. Geurts 1998) use 'denial' to pick out these assertion-challenges. That's, of course, not the kind of denial I'm focusing on.

10.2.5 Reason for hope

It does not seem, then, that we can have a theory of denial parallel to our theory of assertion unless we include the D operator in our language. And, as we've seen, D, in the presence of D-exhaustivity and D-exclusivity, causes problematic paradox. But all is not lost—we have reason to question D-exclusivity and D-exhaustivity (or at least their conjunction) anyway. This reason is provided in Restall (2013): we have problematic paradox from these two principles alone, even without D!

Restall presents constraints on assertion and denial in a sequent calculus with a particular interpretation (motivated in Restall (2005)). Where Γ and Δ are sets of sentences, Restall reads $\Gamma \vdash \Delta$ as: it is incoherent to assert everything in Γ and deny everything in Δ .

Given this reading of \vdash , D-exclusivity is easy to express:

$$(\text{Id}) \quad \Gamma, A \vdash A, \Delta$$

That is, no matter what else we assert (Γ) or deny (Δ), it's incoherent to both assert and deny A. D-exhaustivity is a bit trickier; however, it justifies:

$$(\text{Cut}) \quad \frac{\Gamma, A \vdash \Delta \quad \Gamma \vdash A, \Delta}{\Gamma \vdash \Delta}$$

That is, if asserting A is incoherent given your other commitments, and denying A is incoherent given your other commitments, then your other commitments are already incoherent on their own.

From (Id), (Cut), and a number of principles of naive set theory (or naive truth), Restall shows how to derive $\vdash p$. That is, for any p, it's incoherent to deny p. We could question the principles of naive set theory or naive truth, of course, but then we wouldn't be playing the dialetheist/paracomplettist game we set out to play. So it looks like either D-exclusivity or D-exhaustivity has to go, even with no D in the language.

This is *good* news. After all, it was only in the presence of D-exclusivity and D-exhaustivity that D caused any trouble. Perhaps, then, we can add D to our language while avoiding the trouble, if we figure out just how to weaken these principles.

10.3 Negation and Denial

10.3.1 \neg and D side by side

In this section, I briefly consider ways to add a unary operator D to the language to express denial. So far, I've considered both dialetheist and paracomplete theories

of negation. For concreteness here, I'll assume negation works as dialetheists suppose (that is, that it doesn't satisfy $A \wedge \neg A \vdash \perp$). If we were to instead look at paraconsistent theories of negation (where $\top \vdash A \vee \neg A$ fails), this would all play the same, *mutatis mutandis*.

We can't have both (D-EFQ) $A \wedge DA \vdash \perp$ and (D-LEM) $\top \vdash A \vee DA$. Remember the denier— $\delta = DT\langle\delta\rangle$, where $T\langle\rangle$ is intersubstitutable, and consider the following argument:

| | | |
|----|---|------------------------|
| 1 | \top | |
| 2 | $\delta \vee D\delta$ | 1, D-LEM |
| 3 | δ | |
| 4 | $T\langle\delta\rangle$ | 3, Capture |
| 5 | $DT\langle\delta\rangle$ | 3, Substitution |
| 6 | $T\langle\delta\rangle \wedge DT\langle\delta\rangle$ | 4, 5, \wedge -I |
| 7 | \perp | 6, D-ECQ |
| 8 | $D\delta$ | |
| 9 | $DT\langle\delta\rangle$ | 8, Intersubstitutivity |
| 10 | $DDT\langle\delta\rangle$ | 8, Substitution |
| 11 | $DT\langle\delta\rangle \wedge DDT\langle\delta\rangle$ | 9, 10, \wedge -I |
| 12 | \perp | 11, D-ECQ |
| 13 | \perp | 3–7, 8–12, \vee -E |

10.3.2 Gappy D

Here, I consider relaxing D-LEM, while retaining D-ECQ. Can this get us a denial fit to express disagreement? We've got half of what we want: if I assert DA , I'd better disagree with A (on pain of explosion). Unfortunately, there will be sentences I disagree with, but will not assert the D-sentence of.¹³

For example, consider the denier δ . As the above argument shows, $\delta \vee D\delta$ must fail, and so I should disagree with anyone who asserts δ , and disagree with anyone who asserts $D\delta$. But how? If I disagree with a δ -asserter by asserting $D\delta$, I'm in

¹³ The situation for D here is exactly the situation for the so-called 'arrow-falsum' connective used by Priest and Beall to force denial in some cases.

trouble. I want to disagree with the $D\delta$ -asserter too, but now I *am* a $D\delta$ -asserter. So something's gone wrong.

Negation is no more help here than it ever was. I might assert $\neg\delta$ and $\neg D\delta$ (in fact, I have to, if $\top \vdash A \vee \neg A$ holds), but this isn't enough to express disagreement—I assert $\neg\lambda$ and $\neg D\lambda$ as well, but I don't disagree with a λ -asserter. (Remember, I'm here assuming the dialetheist is right about negation, and in particular about the liar λ .)

If I need to invoke a new speech act, say of *shmenial*, to express my disagreement with the denier, then something's gone very wrong. So it doesn't look like I have the resources to express disagreement with a gappy D .

10.3.3 Glutty D

On the other hand, if D is glutty (that is, if we relax D -ECQ instead of D -LEM), then it's not clear how D helps to express disagreement any more than we already could with negation alone. We should assert $D\delta$, but we don't disagree with δ . The situation is just the same as with negation and the liar.

We might try taking both \neg and D to be glutty, require both LEM and D -LEM, and add a new requirement: $A \wedge \neg A \wedge DA \vdash \perp$. In this system, negation can't express disagreement on its own, and neither can denial. But together, they suffice to force disagreement. Is this any better?

Actually, it's worse; the system is trivial. Let $\iota = \neg\iota \wedge D\iota$. (The proof relies on distribution to show $\top \vdash A \vee (\neg A \wedge DA)$. From there, it's familiar.)

10.4 Making Do

To sum up so far: classicalists can use negation to express disagreement, but para-completists and dialetheists have trouble following suit. To express disagreement, they appeal to denial as a separate sort of speech act, one that satisfies D -exclusivity and D -exhaustivity, and so can be used to express disagreement where negation fails to work. The arguments in section 10.2 have tried to show that giving a complete theory of denial requires us to have some connective D in the language that embeds denial, in the sense that asserting DA is equivalent to denying A , for every A . I've argued that one of D -exclusivity and D -exhaustivity must fail, if we are to have D without triviality. In section 10.3, we saw that, without D -exclusivity and D -exhaustivity, this D is going to have the same trouble expressing disagreement that negation originally had. We seem to have gotten nowhere.¹⁴

¹⁴ To summarize the summary: we want a connective D to embed denial, but any such connective will fail to express disagreement for the same reasons negation failed originally.

This suggests that the strategy of appealing to denial as a separate speech act, and distinguishing D from negation, was pointless in the first place. The para-completist and dialetheist alike should accept, with the classicalist, that negation embeds denial,¹⁵ and that asserting the negation of A is equivalent to denying A. There is thus no need to distinguish denial from assertion of negation, and our theory of denial can become part of our theory of assertion. (Alternately, we could maintain the distinction between denial and assertion of negation, and simply use negation as section 10.2 suggests we use D to construct theories of denial from our theories of assertion.)

We have one remaining choice to make: we can continue to accept the arguments that negation can't express disagreement, and so accept that denial doesn't express disagreement, or we can find a flaw in those original arguments, and keep negation, denial, and disagreement all tied together. I'll explore one variety of each option in turn, although there may well be others.

10.4.1 Exclusion

Suppose we accept the arguments that negation can't express disagreement. If we suppose that negation embeds denial, then denial doesn't express disagreement either. How can we understand disagreement on such a view?

One possibility is to invoke a binary relation on the propositions themselves, and say that two people disagree when there is a pair x, y of propositions such that one asserts x , the other asserts y , and x and y are related. Following Marques (2014), I'll call this relation *exclusion*. For our purposes here, we can take exclusion to be primitive: some things simply exclude other things. For example, we can assume that that the Cheshire Cat is on fire excludes that the Cheshire Cat is in a lake. Then if Alice asserts that the Cheshire Cat is on fire, Humpty Dumpty can indicate his disagreement by asserting that the Cheshire Cat is in a lake. No appeal to denial or negation is necessary to disagree on this account.¹⁶

Something like this might work, but there is familiar trouble lurking not far away. Consider the *excluder*: $\eta = \text{'}\eta \text{ is excluded by some true content'}$. Well, if η

¹⁵ That negation embeds denial is argued for, on very different grounds, in Price (1990).

¹⁶ Marques's discussion takes careful note of possible shifts in context and circumstance of utterance, which I am ignoring here. In some ways, this notion of exclusion is quite like the relation of *incompatibility* drawn on in discussions of negation by Dunn (1993); Restall (1999). But that relation isn't between contents or propositions. The discussion of negation in Brady (2006, 20–1) is also related, as is Millikan (1984, 224–9).

All of these discussions tie exclusion and its relatives very closely to negation, which is no good for the purposes I'm exploring here. I'm assuming that the tie can be severed without too much loss. But since I'm about to argue that this strategy won't work anyway, that's not really a worrisome assumption.

is true, then it's excluded by some true content. But it at least seems that that last sentence expresses a true content that excludes η . So η is excluded by some true content, and so it's true. We should assert η and a content that excludes η . So it looks like exclusion isn't sufficient for disagreement.

Perhaps there's a way to give an account of exclusion, or something like it, that gets around this. But I think there is a more natural account available anyway, and it's to that that I now turn.

10.4.2 *Paracoherentism*

The more natural account takes negation, denial, and disagreement all to be tied together, just as the classicalist thinks. That is, on this account, negation embeds denial, and denial expresses disagreement. Thus, the dialetheist, in asserting $\lambda \wedge \neg\lambda$, disagrees with herself. This amounts to a different response to the original problem of disagreement. On this approach, the dialetheist picture is *incoherent*: it asserts and denies the same thing, and it takes assertion and denial to express disagreement. The question for the dialetheist thus becomes: how bad is incoherence?¹⁷

If incoherence is to be anything less than crippling, it had better be possible to be incoherent in a limited way. That is, we must be able to have an incoherent take on a content A without necessarily having an incoherent take on any old content B . If we replace 'incoherent' here with 'inconsistent', though, this is a familiar problem, and it is to be solved by adopting an appropriate paraconsistent logic, in which one can be inconsistent about A without having to be inconsistent about B . In the present setting inconsistency amounts to incoherence, and thus a paraconsistent approach amounts to a *paracoherent* approach, in which one can be locally incoherent without global incoherence. Since we know there are appropriate paraconsistent logics, there should be no trouble here. Let's christen this sort of view—dialetheism plus the view that negation embeds denial and denial expresses disagreement—*paracoherentism*.

Taking the paracoherentist option I'm suggesting here amounts to choosing to maintain a certain amount of *reflective tension*. This is not an equilibrium position; someone who adopts it disagrees with themselves, after all. On this view, that is the only way to believe truly. Although the truth cannot be coherently stated or believed, it still can be both stated and believed. It simply requires cultivating the right sort of non-equilibrium state.

¹⁷ I focus on dialetheism in this section; as usual, the paracomplete approach plays largely the same. For example, on this picture, the paracompletist, in refusing to assert $\lambda \vee \neg\lambda$, disagrees with nobody, not even the dialetheist. The parallel question is: how bad is it to fail to disagree with one's opponents? How close is it to not having a view at all?

10.4.2.1 ADVANTAGES TO PARACOHHERENTISM

Paracoherentism has several nice features. For one thing, it allows us to answer the demands of section 10.2 by providing an operation on content that embeds denial. In particular, we can 1) state norms on denial parallel to our norms on assertion—(Deny-F) and (Deny-KF) will do as parallels to (Assert-T) and (Assert-K), respectively; 2) express agreement using truth; and 3) explore the relation between denial and falsity as parallel to the relation between assertion and truth. We can define D as suggested from the Stalnakerian framework, and we see that D just is \neg ; our theory already includes it.

What's more, paracoherentism allows us to use denial to express disagreement; we don't need a separate theory of disagreement, as we would on the approach explored in §10.4.1. And we solve the initial disagreement puzzle.

10.4.2.2 CHALLENGES FOR PARACOHHERENTISM

Many of the natural objections to paracoherentism have natural analogs as objections to dialetheism. For example, Slater (1995) objects to dialetheism on the grounds that if $A \wedge \neg A$ can be true, \neg must not be a real negation. One could similarly object to paracoherentism by claiming that if one can rationally disagree with oneself, 'disagree' must not pick out real disagreement. Priest (2006a) responds to Slater by pointing out how many of negation's features his \neg has. Not least among these is the preservation of $\neg(A \wedge \neg A)$ as a theorem-scheme. Similarly, we can respond to the corresponding objection by pointing out how many of disagreement's features disagreement retains on this theory. Not least among these is Dis-exclusivity:

Dis-exclusivity: Agreement and disagreement are *incompatible* states

We can (and should) hold to this. Agreement and disagreement really are incompatible; it's incoherent to do both.

Now, it's possible to find this reply unconvincing. But then, I suggest, one ought to find Priest's original reply to Slater unconvincing as well. *For the dialetheist*, paracoherentism is no extra cost on these grounds. The argument purporting to show that incoherence is bad mirrors the argument purporting to show that inconsistency is bad. If one is unconvinced by the latter, one ought to be unconvinced by the former as well.

A novel problem for paracoherentism (that is, a challenge faced by paracoherentists but not by dialetheists) is in giving an account of logical consequence. For example, suppose we adopt the suggestion in Restall (2005) of taking $\Gamma \vdash \Delta$ to indicate that it's incoherent to accept all of Γ and reject all of Δ . Then, on the present

view, $A \wedge \neg A \vdash B$; since it's incoherent to accept $A \wedge \neg A$, it's incoherent to accept $A \wedge \neg A$ and reject B . Thus, a paracoherentist view can't be closed under \vdash , on this understanding of \vdash ; this is not a paracoherent relation.

Alternately, we might follow Beall and Restall (2006) in taking $\Gamma \vdash \Delta$ to indicate that there is no case in which everything in Γ is true and everything in Δ is not true. This would require us to give some theory of truth-in-a-case. In particular, we would have to give a theory about when a negation is true-in-a-case. And we should be careful. If there is no case where $A \wedge \neg A$ holds, then $A \wedge \neg A \vdash B$ will hold. A paracoherentist view can't be closed under this reading of \vdash either. This is so even if there are some cases where $A \wedge \neg A$ holds and B fails to hold; in this case, $A \wedge \neg A \vdash B$ would fail as well. Priest (2006a) suggests something like this, but phrases the definition slightly differently (using a restricted universal quantification: every case at which everything in Γ is true is also such that something in Δ is true). He does not get the bad result, but it is avoided only by appealing to a non-classical metalanguage that, as yet, awaits full development. The paracoherentist might pursue this general approach, however; she just needs to be a bit careful.

A final possibility for understanding consequence follows Brady (2006) in taking $\Gamma \vdash \Delta$ to express a *containment* between the contents of Γ and Δ . Brady takes this quite literally, using sets of sentences as contents and taking containment to be ordinary set-theoretic containment. On his view, $\Gamma \vdash \Delta$ whenever the union of the contents of the Γ s, closed in a certain way, contains the intersection of the contents of the Δ s. As far as I yet see, this route holds no pitfalls for the paracoherentist, but further exploration will have to wait for another day.

10.5 Conclusion

The initial disagreement problem was supposed to show that negation can't express disagreement for a dialetheist or paracompletist. By and large, dialetheists and paracompletists have accepted this argument and appealed to a speech act of denial, separate from assertion and negation, to express disagreement. They've taken denial to be parallel to assertion. However, when we follow that through, we see that there is real trouble in taking the parallel seriously unless there is an operation D on content such that an assertion of DA is equivalent to a denial of A . What's more, some ways of understanding denial give us the resources to define such a D . This seems problematic, since we can form a new paradox (the denier) with D that dialetheism and paracompletism alone don't address.

In trying to address this paradox, we've seen that we end up saying the same things about D that the dialetheist and paracompletist already said about negation.

Thus, we end up facing the same puzzle about using D to express disagreement as they already faced about using negation to express disagreement. This suggests that distinguishing \neg from D in the first place was a mistake. Negation embeds denial. This, of course, leaves us with our initial puzzle about disagreement intact.

We can try to solve the puzzle by appealing to something other than negation/denial to express disagreement, or we can attempt to keep the three closely linked. I recommend the latter course. The dialetheist approach, seen in this light, is incoherent, but only locally so. It is possible to be sensibly incoherent. The approach has several advantages over orthodox dialetheism, and seems to face few new troubles not also faced by dialetheism. One new trouble—over understanding logical consequence—may well be solvable.¹⁸

References

- Beall, Jc. 2009. *Spandrels of Truth*. Oxford: Oxford University Press.
- Beall, Jc and Restall, Greg. 2006. *Logical Pluralism*. Oxford: Oxford University Press.
- Brady, Ross. 2006. *Universal Logic*. Stanford, CA: CSLI Publications.
- Dunn, J. Michael. 1993. "Star and perp: two treatments of negation." *Philosophical Perspectives* 7:331–57.
- Field, Hartry. 2008. *Saving Truth from Paradox*. Oxford: Oxford University Press.
- Geurts, Bart. 1998. "The mechanisms of denial." *Language* 74:274–307.
- Horn, Laurence R. 1985. "Metalinguistic negation and pragmatic ambiguity." *Language* 61:121–74.
- Horn, Laurence R. 1989. *A Natural History of Negation*. Chicago, IL: University of Chicago Press.
- Marques, Teresa. 2014. "Doxastic disagreement." *Erkenntnis* 79:121–42.
- Millikan, Ruth Garrett. 1984. *Language, Thought, and Other Biological Categories*. Cambridge, MA: MIT Press.
- Parsons, Terence. 1984. "Assertion, denial, and the liar paradox." *Journal of Philosophical Logic* 13:137–52.
- Price, Huw. 1990. "Why 'not'?" *Mind* 99:221–38.
- Priest, Graham. 2006a. *Doubt Truth to be a Liar*. Oxford: Oxford University Press.
- Priest, Graham. 2006b. *In Contradiction*. Oxford: Oxford University Press, 2nd edition.

¹⁸ For discussion regarding earlier versions of this chapter, thanks to audiences at the University of Melbourne, the Arché Logic of Denial workshop, the Propositional Content and Proposition-Related Acts workshop, and PALMYR IX: Logic and the Use of Language. Comments from Aaron Cotnoir, Catarina Dutilh Novaes, Graham Priest, Greg Restall, and Stewart Shapiro proved especially helpful at these events. This research was partially supported by the Agence Nationale de la Recherche, program 'Cognitive Origins of Vagueness', grant ANR-07-JCJC-0070, by the project 'Borderlineness and Tolerance' (Ministerio de Ciencia e Innovación, Government of Spain, FFI2010-16984), and by the Australian Research Council Discovery Project 'Paraconsistent Foundations of Mathematics'; I'm grateful to all of them.

- Restall, Greg. 1999. "Negation in relevant logics: how I stopped worrying and learned to love the Routley star." In D. Gabbay and H. Wansing (eds.), *What is Negation?* 53–76. Dordrecht: Kluwer.
- Restall, Greg. 2005. "Multiple conclusions." In P. Hajek, L. Valdez-Villanueva, and D. Westerstahl (eds.), *Proceedings of the Twelfth International Congress on Logic, Methodology and Philosophy of Science*, 189–205. London: King's College Publications.
- Restall, Greg. 2013. "Assertion, denial, and non-classical theories." In Koji Tanaka, Francesco Berto, Edwin Mares, and Francesco Paoli (eds.), *Paraconsistency: Logic and Applications*, 81–100. Dordrecht: Springer.
- Slater, Hartley. 1995. "Paraconsistent logics?" *Journal of Philosophical Logic* 24:451–4.
- Stalnaker, Robert. 1978. "Assertion." *Syntax and Semantics* 9:315–32.