

A user's guide to hybrid tools

One of the main innovations in recent metaethics has been to take moral judgments to be *hybrid* states, consisting of more than one component. On one particularly prominent version, moral judgments are hybrids of a desire-like component and belief-like component. This hybrid idea promise to cut the Gordian knots that have animated metaethical debate. Moral judgments have some features that are naturally explained by taking them to be desire-like: for example, they seem linked to motivation. But moral judgments have other features that are naturally explained by denying that they are desire-like. Hybrid theories can capture all those features by appeal to the different components that they take to constitute a moral judgment. One component captures the desire-like features of moral judgment, and the other the features that don't seem desire-like.

Unfortunately, though, hybrid theories face very serious problems. The problems are *so* serious that they seem to show that no philosophically interesting hybrid theory could be correct. One central problem is that philosophically interesting hybrid theories seem to force the wrong predictions about attitude ascriptions. I'll explain how to disarm this central problem. My solution is to take it to be an instance of a highly general problem: a problem about what are sometimes called 'intensional anaphora', and to show that *any* adequate account of intensional anaphora solves the problem.

After introducing my solution, I'll also suggest that the solution is highly *fruitful*. It solves a range of central problems for hybrid theorists. Now there won't be space to fully show that it solves all those other problems. But my overarching goal will to raise the bar for objections to hybrid views. I'll claim that an objection is compelling only if it applies to hybrid views that are developed in the way I suggest. And, as we'll see, many initially compelling objections to hybrid views lose much of their force given my suggestion.

Now many metaethicists are suspicious of hybrid tools. They see them as one of the least legitimate elements in our toolkit. My ambition is to flip this impression: to show that hybrid tools are among the *most* legitimate tools in our toolkit. In doing that, I'll

focus on the use of hybrid tools in metaethics. But metaethics has been a rich source of resources for other subfields, with ideas like expressivism and notions like supervenience taken up broadly after explicit discussion in metaethics.¹ Hybrid tools are similarly ripe for cross-(sub)-disciplinary appropriation. So it's not just metaethicists who should care about the legitimacy of hybrid tools. If hybrid tools can do interesting metaethical work, they're available for interesting applications in other subfields, too.

1 The promise of hybrid tools

I've introduced hybrid theories as accounts of 'moral judgment', using 'moral judgment' to refer to whatever mental states are expressed by sincere uses of sentences like 'abortion is permissible'. We can say that traditional cognitivists take moral judgments to be belief-like states, and traditional expressivists take them to be desire-like states.² Hybrid theorists take moral judgment to consist in more than one state – for example, in both a belief and a desire-like state. Some hold that attitude ascriptions *semantically express* both states. Others hold that only one state is semantically expressed, while the other is either conventionally or pragmatically communicated.

But however the details go, hybrid theories look like promising ways to capture at least four different kinds of philosophical advantages. First, hybrid tools promise to explain connections between moral judgment and motivation. It is hard to see how a rational person could sincerely judge that eating meat is wrong unless she's motivated to not eat meat. And this connection seems like evidence

¹Sarah Moss (2017), Paolo Santorio (2016), and Seth Yalcin (2007) have recently drawn on expressivist tools for applications outside of metaethics. Hare was the first to use the word *supervenience* in its modern sense (Hare 1952, p. 80,131), and, as Selim Berker (2017) noted, the same idea appears in (Moore 1922, p. 261), Henry Sidgwick (Sidgwick 1907, p. 209, 379), and W. D. Ross (Ross 1930, p. 109, 120, 122–23)

²I myself think that the central disagreement is about the nature of moral *belief*. If I sincerely judge that abortion is permissible, I'm also ready to assert sentences like 'I believe that abortion is permissible'. And if I'm ready to assert those sentences, I'll also be ready to semantically descend, and say that I believe that abortion is permissible. But it's more familiar to formulate the debate as a debate about the nature of moral judgment, so that's what I'll do here. For further discussion, see Jamie Dreier (2005), on what he calls 'the problem of creeping minimalism'.

that moral judgment doesn't consist in a belief, at least if we think that belief is itself motivationally inert.

Now cognitivists can and do offer explanations of moral motivation. Some posit background motivational states that combine with our moral beliefs to explain ordinary cases of moral motivation. But many philosophers find these extensions unsatisfactory. They grant that the explanations may explain ordinary cases of moral motivation. But they deny that the explanations capture the *intimacy* of the connection between moral judgment and motivation.³ Other cognitivists insist that beliefs can motivate by themselves, without desire-like states, and thereby find themselves entangled in highly fraught debates in moral psychology.⁴

Hybrid theorists can do better than traditional cognitivists. Suppose that moral judgments consist of *both* beliefs and desire-like states. If the desire-like state is semantically expressed, or is conventionally implicated, the hybrid theorist can capture the intimacy of the connection between moral judgment and motivation just as well as expressivists. The agent is *guaranteed* to have the desire-like state if the moral judgment is sincere, since the desire is a component of moral judgment. And if the desire-like state is pragmatically communicated, the connection won't quite be that intimate. But it should be *more* intimate than traditional cognitivist predict.⁵

The second advantage of hybrid tools is that they can help explain *inconsistency* between moral judgments. Consider one of the metaethicist's standard punching bags: old-fashioned speaker subjectivism, where my moral judgments ultimately consist in nothing more than beliefs about my own attitudes. Maybe, for example, my judgment that ethnic cleansing is wrong consists of the belief that I dislike ethnic cleansing. As C. L. Stevenson (1937) famously noted, it's hard for this punching-bag view to explain how inconsistent moral judgments really are inconsistent. The punching-bag view takes Hitler's judgment that ethnic cleansing is obligatory to be the belief that Hitler likes ethnic cleansing. So it seems to predict that Hitler's judgment is consistent with my judgment. After all, it's perfectly coherent to believe that Hitler likes ethnic cleansing while believing that I don't.

³Simon Blackburn puts this complaint with particular force (1998, 84ff).

⁴Michael Smith (1994) has a classic discussion of the issues here.

⁵See Stephen Barker (2000), Daniel Boisvert (2008), David Copp (2001), and Caj Strandberg (2012, 2015).

Hybrid tools can help explain inconsistency between moral judgments. As one possible toy illustration of how hybrid tools can help, consider a version of the punching-bag view that is extended with a further *intention*. On this suggestion, my judgment that ethnic cleansing is wrong would be a hybrid of the belief that I dislike ethnic cleansing plus the intention not to do what I dislike, and that Hitler’s judgment is a hybrid of the belief that Hitler likes ethnic cleansing plus an intention to do what Hitler likes. And assume that intentions are intentions to perform actions, rather than to bring about states of affairs. Hitler’s intention is an intention to do actions that Hitler likes, rather than the intention to bring it about that Hitler does actions that Hitler likes.⁶

The resulting view can explain why no one could coherently have my judgment and Hitler’s judgment. If you have Hitler’s moral judgment, you intend to do what Hitler likes, which includes committing ethnic cleansing. So you intend to commit ethnic cleansing – that is, you intend for yourself to commit ethnic cleansing. And if you have my judgment, you intend for yourself not to commit ethnic cleansing. And no one could coherently have both intentions, since you can’t consistently intend for yourself to do inconsistent things. That’s why no one could coherently accept my moral judgment and Hitler’s moral judgment.⁷ Now, as we’ll see later, this suggestion faces further problems. But it at least can explain why inconsistent moral judgments really are inconsistent.

Hybrid tools could also be used in moral epistemology. Consider *global* moral skeptics who argue that we don’t have any substantive moral knowledge. Such skeptics argue both for a conceptual claim, about what would be required for moral knowledge, and a substantive claim, about propositions that we don’t know. Hybrid tools could be used to resist the conceptual claims that the skeptic relies on.

Here is one concrete illustration. Suppose that in asserting ‘killing is usually wrong’, I’m committing myself to two propositions: a proposition expressing the natural fact that some particular set of rules usually forbids killing, and a proposition expressing the moral fact that that very same set of rules captures the demands of moral-

⁶Mark Schroeder (2011) defends this conception of intention. I’m not assuming that this conception of intention is correct. I’m just using it as an illustration.

⁷Steven Finlay (2014) and Alex Silk (2016) favor views that share some problems about inconsistent moral judgments with old-fashioned speaker subjectivism. And they have both drawn on hybrid tools to answer those problems.

ity. A hybrid theorist could analyze moral knowledge as a hybrid of knowledge of the natural fact about how that set of rules classifies actions, plus mere acceptance of the moral fact that that set captures the demands of morality. Then moral knowledge could be easier to acquire, since the norms on acceptance could be less strict than the norms on knowledge.⁸ Then it would be easier to answer moral skeptics, because their conceptual claims about moral knowledge would be mistaken. The mistaken conceptual claim would be that knowledge of natural facts about how a set of rules classifies actions isn't the sort of knowledge required for moral knowledge – that moral knowledge requires *knowledge* rather than mere *acceptance* of propositions expressing the moral facts. This is the application of hybrid tools that most interests me; see Perl (2017, forthcoming).

And the fourth possible application of hybrid tools is by *expressivists* to answer the Frege-Geach problem, to explain the inferences that our moral judgments license. Hybrid expressivists extend a basic expressivist theory by adding further descriptive beliefs, and using those further descriptive beliefs to capture the inferential relations between moral judgments.⁹

2 A faultline in hybrid theories

But when we turn to critically evaluating hybrid tools, we find serious problems. In fact, the problems are *so* serious that we can reasonably doubt if it's possible to use hybrid tools.

Let's begin by noting that hybrid theorists need to associate sentences S and $\lceil \neg S \rceil$ with mental states that no one person could consistently have at the same time. Doing that seems like a necessary if not sufficient condition for explaining cases of moral disagreement.

BASIC INCONSISTENCY CHALLENGE: associate S and $\lceil \neg S \rceil$
with mental states that no one could consistently have.

In this paper, I'll assume that hybrid theorists do somehow address this CHALLENGE. We saw earlier how hybrid ideas can make the

⁸I'm here using 'acceptance' as a placeholder for an attitude that is not normed by knowledge – for one possible account of such an attitude, see Stalnaker (1998).

⁹For discussion, see David Alm (2000), Dorit Bar-On and Matthew Chrisman (2009), Frank Jackson (1999), Richard Joyce (2006), and Michael Ridge (2006, 2014).

CHALLENGE easier to answer. The punching-bag view, old-fashioned speaker subjectivism, struggled to answer this CHALLENGE. But a hybrid addition (positing a further intention) helps. It's reasonable to think that many different kinds of hybrid views will have the resources to answer this CHALLENGE.

The rest of the paper will focus on the serious and possibly devastating problems that arise after the hybrid theorist has answered the CHALLENGE. To introduce the problems, we need to start with an important faultline among different hybrid theories. Consider the hybrid theories that explain moral motivation by taking moral judgments to be a hybrid of a belief and a desire. Let S be a moral sentence that's free of indexicals – the sentence *abortion is permissible*, say.¹⁰ One hybrid view is that *anyone's* judging that S *always* consist in the very same belief and the very same desire. I'll call this kind of hybrid view a *cautious* hybrid view. One cautious hybrid view is that *anyone's* judging that abortion is permissible *always* consists of the belief that abortion is compatible with what an impartial spectator approves and a desire to act as the impartial spectator approves.¹¹

An alternative kind of hybrid theory allows that judging that S can consist in different belief/desire pairs for different people. Imagine that Jeremy has a desire to maximize happiness, and David has a desire to act in a way that the impartial spectator approves. This second kind of hybrid theory would take Jeremy's moral judgments to consist in a belief about happiness plus a desire about happiness, and take *David's* moral judgments to consist in a belief about the impartial spectator plus a desire about the impartial spectator.

I will call this second kind of hybrid theory an *ambitious* hybrid theory. Why call them *ambitious*? Traditional accounts of the assertive use of 'abortion is permissible' have two features. First, they take that use to express just one mental state – either a belief-like state, or a desire-like state. Second, they take that use to express the *same* mental state for all speakers. Cautious and ambitious hybrid theorists all abandon the first feature of traditional accounts. But ambitious accounts also abandon the second feature. I call them ambitious because they involve a more significant departure from

¹⁰I throughout italicize sentences I intend to mention.

¹¹It's standard to present hybrid views in the way that I've just done, as involving, say, beliefs about impartial spectators. I take it that this presentation is a harmless abstraction from the particular details of the hybrid view.

tradition.

Extant hybrid theories have tended to be ambitious rather than cautious.¹² And the reason is that cautious hybrid theories have comparatively little philosophical payoff. Consider the cautious explanation of moral motivation. It posits belief/desire pairs that we all share: say, belief/desire pairs about what an impartial spectator approves. And this posit is first semantically implausible. It is difficult to find some such belief that all uses of ‘abortion is permissible’ express merely by virtue of the sentence’s meaning. Second, this posit assumes a striking uniformity in human psychology. And this striking psychological posit again seems implausible. Even worse, this striking psychological posit already explains by itself much of what needs to be explained. Mark Schroeder puts the point this way.

we can construct an ordinary externalist, cognitivist theory, which assigns moral sentences the same descriptive contents as the hybrid view does and makes the same background assumption about desires but does not claim that these desires are expressed by the sentences. ...Then whenever anyone forms a new moral belief about what is wrong, the externalist cognitivist imitator explains that she will be motivated by this background desire not to do that thing. (Schroeder 2009, p. 302)

The resulting view looks like cognitivism plus motivational externalism. And the use of hybrid tools doesn’t help make the explanation more plausible; you can dispense with them and not lose anything. A similar point holds for other dramatic philosophical payoffs. Now there are some other philosophical payoffs that hybrid tools cautiously used can secure.¹³ But they’re not the payoffs advertised in §1, and it’s much more controversial if they’re payoffs that are worth securing.

¹²Daniel Boisvert (2008) has been the main exception.

¹³Mark Schroeder gives some examples (Schroeder 2009, p. 301, 303). For example, he notes that they can explain the intuition that there is a conceptual and necessary connection between moral judgments and motivation. But he also suggests that the connection that they vindicate is too strong to be plausible.

3 The challenge to ambitious hybrid theories

I've just described an important faultline among hybrid theories, and noted that the hybrid theorist has to be on the ambitious side of the faultline if her hybrid tools are to help secure the striking advantages described in §1.

Unfortunately, though, ambitious hybrid theories face serious problems of their own, problems which appear to show that no ambitious hybrid theory could be correct. One basic problem is about attitude reports. Consider the ambitious hybrid theory that associates Jeremy's moral judgments with beliefs and desires about maximize happiness, and David's moral judgments with beliefs and desires about impartial spectators. Jeremy and David could both agree on some moral question – for example, they could both agree that abortion is permissible. And they could be perfectly rational in agreeing on that point, if there are cases where abortions are compatible with what maximizes happiness and also cases where an impartial spectator would approve of an abortion.

So this hybrid theory needs to explain how sentences like (1) could be used to assert something true.

- (1) Jeremy and David agree that abortion is permissible.

The easiest explanation would be that the use of (1) asserts that

- Jeremy agrees that abortion is compatible with what maximizes happiness and desires to maximize happiness, and David agrees that abortion is compatible with what an impartial spectator approves, and desires to do what an impartial spectator approves.

But this explanation cannot work, because it violates a general stricture about 'agrees'. The general stricture is that terms in the complement of an 'agrees' report have to make the same contribution to what is asserted about each agent. This explanation violates this stricture by taking 'permissible' to make a different contribution to what is asserted about Jeremy than to what's asserted about David. To motivate the general stricture, contrast (a) and (b).

- (a) Jeremy and David *each believe* that a local newspaper just shut down.

(b) Jeremy and David *agree* that a local newspaper just shut down.

(a) can be used to assert that Jeremy and David have beliefs about different newspapers that are local to each of them – one in New York, and one in LA, for example. On this use, ‘local’ is interpreted differently for David than it is for Jeremy. But (b) cannot be used in that way, where ‘local’ is interpreted differently for David than it is for Jeremy. ‘Local’ has to make the same contribution for both agents in ‘agrees’ reports.¹⁴

So whatever a use of (1) asserts, it has to take ‘permissible’ to make the same contribution for both agents. As a result, there are only three kinds of options for what a use of (1) could assert.

- Jeremy and David agree that abortion is compatible with what *maximizes happiness*, and desire to maximize happiness.
- Jeremy and David agree that abortion is compatible with what *an impartial spectator would approve*, and desire to do what an impartial spectator would approve.
- Jeremy and David agree that abortion is compatible with what is P, and desire to do what is P – where P is a property neither about happiness nor about impartial spectators.

¹⁴Cappelen and Hawthorne emphasize this general point about ‘agrees’ and use it in evaluating epistemic contextualism (2009, 54ff).

There are some complications about this data. Consider (c):

(c) Given Jeremy’s experiences in New York and David’s experiences in LA, they agree because of independent evidence that local newspapers are under threat.

If we focus enough on the material about New York and LA, we might interpret Jeremy’s evidence as about New York, and David’s evidence as about LA. So (c)’s appropriateness might seem to show that ‘local’ makes a different contribution for each agent. But note that I can say discourse-initially, ‘local newspapers are under threat’. In so saying, I am not just talking about the newspaper(s) local to me. In this case, we should posit an implicit, unselective quantifier like *usually* or *always* that binds the location variable, so that the sentence communicates that usually (newspapers local to x are under threat) Lewis (1975). We should say the same thing about (c). But then (c) is compatible with the generalization in the main text. On the unselective-quantifier interpretation, ‘local’ is *not* interpreted differently for David than it is for Jeremy; for both of them, its interpretation is given by combining its semantic value with the implicit unselective quantifier.

And (1) can be used to assert something true even while each of these options is false. Imagine that Jeremy believes that an impartial spectator would disapprove of abortion, and that David believes that abortion isn't compatible with what maximizes happiness. Even while having those beliefs, they can still both judge that abortion is permissible. Jeremy's judgment consists of the belief that abortion is compatible with what maximizes happiness, plus a desire to do what maximizes happiness. And that judgment is fully compatible with his other belief about an impartial spectator. By contrast, David's judgment similarly consists in a belief about impartial spectators. And that judgment is compatible with his belief about what doesn't maximize happiness. In this case, neither of the first two options is true.

In fact, David and Jeremy can share the judgment that abortion is permissible without sharing any particular belief. After all, different belief/desire pairs are sufficient for sharing the judgment that abortion is permissible. So it will be possible for (1) to be true even when the third option is false – when they disagree about what's compatible with what is P, for any P.¹⁵

So the assertability of (1) looks like powerful evidence against this ambitious hybrid view. It's possible to find contexts where (1) can be used to assert something true. And it is hard for the hybrid view to explain how that is possible, given the general constraint on 'agrees'.

This problem is a problem for this particular ambitious theorist. But the same kind of problem will arise for *any* ambitious hybrid theorist whatsoever. Ambitious hybrid theorists hold that judging that S can consist in different mental states for different people. (That's the very thing that distinguishes them from cautious hybrid theorists.) So they allow that two agents can both judge that S without sharing the same mental states. As a result, there will always be some sentence like *abortion is permissible* that two agents can agree on without having the same mental states.¹⁶ And, as we'll see later

¹⁵And if the desire-like state is the only thing that's semantically expressed, we can make the same point with desires – that they could fail to have any relevant desires in common, even while they agree that abortion is permissible.

¹⁶There is one kind of ambitious hybrid view that might evade this problem. Suppose that the moral use of the sentence S always communicates two mental states. One mental state is always the same: S always semantically expresses that mental state. But different mental states are pragmatically communicated by a

on, this problem about agreement is just one illustration of a broader class of problems for the ambitious hybrid theorist.¹⁷

In fact, this kinds of problem is a serious problem for a wide range of metaethical views. It tends to be a problem for *semantically flexible* views where one speaker's use of a sentence like 'abortion is permissible' can express one mental state, while another speaker's use would express a different one. This paper is focusing particularly on ambitious hybrid views because I take them to be the most promising kinds of semantically flexible views – they're the ones best-positioned to answer the BASIC INCONSISTENCY CHALLENGE. But if there are other, non-hybrid ways for semantically flexible views to answer that CHALLENGE, it's worth discussing them with ambitious hybrid theories; they'll face similar problems.

If ambitious hybrid theories can't answer this problem, they're incorrect as descriptive accounts of our moral thought and talk. And ambitious hybrid theories almost all intend their theories as descrip-

use of S – there's always another one, but not always the same one.

The key question for evaluating this kind of view is: is it possible for 'A and B agree that S' to communicate that A accepts an implicature of A's use of S and that B accepts a distinct implicature of B's use of S? If the answer is 'no', the view is immediately implausible; we would 't be able to infer that it would be appropriate for B to assertively utter S if we know that 'A and B agree that S' is true. And that's implausible. Why would that implausible result happen? In order for B's utterance of S to be appropriate, B needs to accept what S would implicate in B's context. And the truth of 'A and B agree that S' doesn't tell us whether he does, given a 'no' answer.

Moreover, it's a tricky empirical question whether the answer to this question is 'yes'. I myself am skeptical. Suppose that in A's context, A's hearer is trying to figure out if Billy can reach the sugar. A says 'Billy is short', implicating that he can't. Suppose that in B's context, B's hearer is trying to figure out if Billy can ride a ride that only allows people under four feet tall. If A says, 'A and B agree that Billy is short', will competent speaker-hearers take A to accept that he can reach the sugar, and B to accept that he can ride the ride? My judgment is that they won't – or at least that it's possible to fill out the context so that they won't. And even if I'm wrong about this case, why wouldn't there be cases that pattern like I take this case to pattern? We need real empirical work on conversational implicatures generated from under 'agrees' to evaluate this suggestion.

¹⁷Mark Schroeder describes these sorts of problems in illuminating detail, in §VI of Schroeder (2009). Brian Weatherson (2008) pointed out the same problem, and Herman Cappelen and John Hawthorne (2009) offer some diagnostic criteria that bring out related problems. The same issues come up in contextualist accounts of epistemic modals, and in contextualist accounts of predicates of personal taste; John MacFarlane (2014) is one of many that forcefully bring those problems out. Though I don't highlight the connections, I take the suggestions in the rest of the paper to solve many of the problems for contextualists as well.

tive accounts, not revisionary accounts.¹⁸ It's thus unsurprising that hybrid theorists have done a great deal of work in tackling this sort of problem. (Approaches include those of Gunnar Björnsson and Steve Finlay (2010), Finlay (2014, 2016), and Michael Ridge (2014).)¹⁹

4 A user's guide for ambitious hybrid theorizing

The rest of this paper will introduce a new approach to these problems. I'll suggest that ambitious hybrid theorists should all adopt my approach, even if they had previously adopted another one.

My approach is to provide a companion in innocence for ambitious hybrid theories. I say that the problem about agreement is just an instance of a highly general problem: a problem about what are sometimes called 'intensional anaphora'. I'll argue that any adequate explanation of intensional anaphora immediately solves the problem about agreement. Since we know that there is some adequate account of intensional anaphora, we know that ambitious hybrid theories are all defensible, at least against this problem.

As a comparison to what I'll do, think of the defense that R. M. Hare (1952) gave against the Frege-Geach problem. He pointed to imperatives as a companion in innocence, and argued that any adequate account of imperatives also solved the problems for his prescriptivism. I'll similarly argue that intensional anaphora are a license for optimism for ambitious hybrid theorists that works just like Hare's license for optimism. It both guarantees that the problems can be solved, and it shows how to solve them. So I won't be critically engaging with other ways of defending or developing of ambitious hybrid theories. If I'm right, those other defenses are trying to solve a problem that our linguistic competence has already solved.

Let's start with (3).

- (3) (a) Bill and Jane agree that a kitten was here. (b)
They also agree that it was happy.

Discourses like (3) are very puzzling. It's not just that Bill and Jane have attitudes about the existential closures of the complements.

¹⁸Though see Richard Joyce (2001) for a possible exception.

¹⁹Others have abandoned ambitious hybrid theories, in favor of cautious ones (Teemu Toppinen (2013) and Mark Schroeder (2013) are perhaps examples).

They definitely agree that some kitten was present. And they definitely agree that something was happy. But that's not all. Their thoughts are linked together; they're not just agreeing that one kitten was here and a distinct one was happy.²⁰

Our question is: what is the correct semantics and pragmatics for (3) that explains how the two attitudes are 'linked' together? The occurrence of 'it' is used anaphorically, to refer to the aforementioned kitten. But the linking is occurring across two different intensional contexts: two attributions of agreement. That's why (3) is said to raise a question about intensional anaphora: it's about a use of *it* that is anaphoric on an expression in a different intensional context.

There are many different accounts of intensional anaphora. My ambition will be to show that each account will provide a solution to the problem from §3 for ambitious hybrid theorists. In fact, I'll argue that any account of intensional anaphora *must* solve that problem if it is to be empirically plausible. Now some accounts of intensional anaphora will support solutions that strongly resemble views that hybrid theorists have already taken. And those solutions will still face the same problems they did when hybrid theorists initially introduced them.

But the core contribution of this paper is to provide a more systematic perspective. For each approach to intensional anaphora, there is a solution to the problem from §3 if that approach is correct. Partisans for other accounts of intensional anaphora may see that solution as indefensible. But they will also think that there is *some other* solution to the hybrid theorist's problem that they regard as defensible – the solution that follows from their favored approach to intensional anaphora.

I'll thus be recommending that ambitious hybrid theorists be less committal than they have been. They should explain how their views would be implemented given each account of intensional anaphora, but be non-committal about which approach to intensional anaphora is correct. That recommendation will be my general solution to the §3 problem. After arguing for the recommendation, I'll turn to some broader questions. I'll first suggest that my solution is available to any hybrid theorist whatsoever. Though my suggestion will involve some rethinking of hybrid views, it won't change any of their philosophically important upshots. Then I'll argue that my solution is

²⁰Peter Geach (1967) drew contemporary attention to these kinds of discourses.

highly *fruitful*: that it has the right structure to solve other important problems. My overarching goal is to raise the bar for objections to hybrid views. An objection is compelling only if it applies to hybrid views that are developed in the way I suggest. And, as we'll see, it's much harder to find such objections.

4.1 The first option: liberal accounts

One natural idea is that the discourse in (3) is about some particular individual. For example, maybe the use of (3) communicates a proposition with an existential quantifier scoped wide.

(5) $[\exists x$ (Bill and Jane agree that x was here, and they agree that x was happy)]

Now open sentences with variables like 'x' paradigmatically contribute singular propositions. So this approach requires Bill and Jane to be thinking singular thoughts about the same object. That is, (3) is true only if there's some particular cat (Miley or Tabby or ...) that they're both thinking about. This account assumes that it is comparatively easy to think singularly about particular objects; it's very liberal in allowing for singular thought. So I'll say that it is the *liberal* account of (3).²¹

Unfortunately, the liberal has to say odd things to capture all the uses of (3). Suppose that Bill thinks that Miley was a kitten who was here and happy, while Jane thinks that Tabby was a kitten who was here and happy. They would then agree that a kitten was here. And they would agree that it was happy. But the liberal can only capture this case by saying – very oddly – that either Bill is thinking of the cat that Jane is thinking about (Tabby), or that Jane is thinking about the cat that Bill is (Miley). There are ways to lessen the oddity.

²¹Many philosophers will find this suggestion controversial, because it requires the ability to think singularly at will. Suppose Bill and Jane know that only kittens are the only animals allowed here and that kittens are always happy when they're here. And they know that one animal was allowed in. So singular truth-conditions for (3) require that purely general grounds can suffice for singular thought. So even philosophers like Robin Jeshion (2010) who are comparatively liberal about singular thought won't accept this account of (3), since she denies that purely general grounds do suffice. But some philosophers are liberal enough about singular thought to accept this account; David Kaplan (1989) and David Manley and John Hawthorne (2012) are all examples. For critical discussion, see (Jeshion 2010, p. 125-9).

Maybe the liberal can suppose that (3)'s truth is grounded in the determinate disjunctive fact that Bill and Jane are both thinking of Tabby, or both thinking of Miley, while insisting that there is no determinate fact about whether they're both thinking about Tabby or both thinking about Miley.

A liberal account of intensional anaphora can solve the §3 problem. Consider the ambitious hybrid theorist who takes moral judgment to be a hybrid of belief and desire. David's attitudes are about what an impartial spectator would approve, and Jeremy's are about what maximizes happiness. They agree that abortion is permissible. The challenge to the ambitious hybrid theorist was to explain what their agreement consist in. Her answer has to have this form:

(AMBITIOUS HYBRID GLOSS) David and Jeremy agree that abortion is compatible with what is K and desire to do what is K.

The problem was that they can agree even without sharing beliefs about what is K.

At this point, we can see that this problem is *exactly the same problem that the liberal faces*.²² So any liberal has to have a solution to it. Maybe, for example, David and Jeremy's agreement is grounded in the determinate disjunctive fact that the AMBITIOUS HYBRID GLOSS is true either of the property *being-what-an-impartial-spectator-approves* or the the property *being-what-maximizes-happiness*, though not in the determinate fact that it's true of the first property or the determinate fact that it's true of the second property. (This suggestion has some resemblances to what someone like Michael Ridge might say (Ridge 2014, p. 118ff).)

I think that this suggestion is odd. But its oddity comes from the liberal account of intensional anaphora, and not from the ambitious hybrid theory. If you accept the liberal account, you should regard this solution to the §3 as perfectly effective.

²²As in this example, some of the hybrid theories are theories about properties, rather than objects. But intensional anaphora can occur with properties as well as objects. Consider *The group agrees that several features of Paris make it more beautiful than Berlin. The group also agrees that they don't make Paris more beautiful than Rome.* Quantifying over features is quantifying over properties, so the agreement in the second sentence is still agreement about a group of properties. I focus on objects in the main text only because the examples are simpler.

4.2 The second option: E-type accounts

I'll now turn to the second kind of leading account of intensional anaphora. This account takes Bill and Jane to agree about a *descriptive* proposition.

(3) (a) Bill and Jane agree that a kitten was here. (b)
They agree that it was happy.

(E-TYPE) Bill and Jane agree that [the x: kitten(x) \wedge x was here] (x was happy).

Gareth Evans (1977, 1980) pioneered this account of anaphora pronouns. It has since been called an 'E-Type' analysis of pronouns.

At first, E-Type accounts look like counterexamples to my claim that any adequate account of intensional anaphora will solve the §3 problem. Consider David and Jeremy, with their different beliefs and desires – one about happiness, and one about impartial spectators. They can agree that abortion is permissible. But it's hard to find a descriptive proposition that they agree on. The proposition needs to have the form '[the P: ...] (abortion is P)'. Options include:

- [the P: P = *permitted-by-the-impartial-spectator*] (abortion is P)

But *Jeremy* need not believe that the impartial spectator would sometimes permit abortion. What matters for his moral judgments is what maximizes happiness, so his agreeing that abortion is permissible doesn't require him to have any beliefs about impartial spectators.

- [the P: P = David desires to do what is P] (abortion is P)

But Jeremy doesn't need to have beliefs about what David desires to believe that abortion is permissible.

- [the P: P = Jeremy and David both desire to do what is P] (abortion is P)

Again, though, agreeing about abortion doesn't require believing that there's *anything* that he and David both desire. It seems like

E-Type theories predict that ambitious hybrid theories *do* face the §3 problem.

Fortunately for hybrid theorists, though, E-Type accounts need to be developed if they are to capture some important further data. And when they are developed, they do solve the §3 problem.

Consider this discourse.

(4) A dog was happy. Bill denies that it was either a dog or happy.

E-Type accounts initially struggle to identify a proposition that Bill can deny. It doesn't seem like he denies the proposition that [the x: dog(x)] (x is a dog or happy). He would be irrational in denying that proposition. And (4) can be true without him being irrational. A similar point holds for the proposition that [the x: dog(x) or happy(x)] (x is a dog or happy). Denying that proposition also would involve a kind of irrationality that (4) need not involve. So in this case, there isn't any informative description that links the proposition that the speaker asserted with the first sentence in (4) with the proposition that Bill is denying.

E-Type accounts thus have to allow that two agents can think descriptively linked propositions even when there isn't any *informative* description that links them together. One way to do that is to take the proposition that Bill denies to be a proposition with a highly general description: something like the proposition that [the x: individual x] (x is either a dog or happy). Since there are many individuals, we need to tacitly ignore those other individuals. A natural way to ignore them is to appeal to highly partial *situations*.²³

SITUATIONIST E-TYPE: For every minimal situation s_1 compatible with Bill's salient beliefs, it's not true in s_1 that [the x: individual x] (x is either a dog or happy).²⁴

If each minimal situation that's relevant only includes one individual, then this suggestion captures the use of (4). And it does that without relying on any informative description that links the two thoughts together.

²³Paul Elbourne (2005) develops this suggestion.

²⁴Since we're glossing a proposition that Bill denies, we make sure that the proposition is not true throughout his salient belief worlds.

Importantly, though, the appeal to situations also solves the problems for the ambitious hybrid theorist.

SITUATIONIST E-TYPE: For every minimal situation s_1 compatible with x 's salient attitudes, it's true in s_1 that [the P: P is the property] (abortion is P) and that x desires to do what is P.

Jeremy and David agree that abortion is permissible is then true iff the SITUATIONIST E-TYPE gloss holds for each of them. What matters is that the minimal situations only include one property – a different one for David than for Jeremy, but only one for each of them.

We should only accept an E-Type account that can explain the use of (4). As a result, the only acceptable E-Type accounts allow that two agents can both think descriptively linked thoughts even when there isn't any informative description that links the thoughts together. But that's just what the ambitious hybrid theorist needs to solve her problems. So any viable E-Type account will also solve the §3 problem for ambitious hybrid theorists.

4.3 The third option: dynamic accounts

I'm in the middle of arguing that any adequate account of intensional anaphora will also solve the problems for ambitious hybrid theorists. So far, I've been arguing by case; I've been showing that extant accounts of intensional anaphora also solve the problems for ambitious hybrid theorists.

At this point, though, you might have started to think that I'm explaining the obscure by the obscure. I'm pointing out that extant accounts of intensional anaphora incur highly unnatural commitments, and then using those highly unnatural commitments to help ambitious hybrid theorists. In leveraging this complaint, you're suggesting that there is something unsatisfying about the §4.1 and §4.2 accounts. (I'm sympathetic with this complaint.)

Fortunately, though, linguists have developed accounts of intensional anaphora that work much better than the ideas in §§4.1–4.2. Their accounts also solve the problems for ambitious hybrid theorists, and in a much cleaner way. If I can appeal to them, I'll explain the obscure by the pellucid.

Their basic idea is to introduce a *sui generis* third way for thinking about objects, which I'll call *thought with unanchored discourse referents*. Thought with unanchored discourse referents is like singular thought, in that it does not require the thinker to think of the object as falling under any particular description. But it's unlike singular thought, in that there need be no unique worldly individual that you're thinking about.²⁵

Unanchored discourse referents model the accumulation of information about objects. They differ from other ways of accumulating information about an object in that they do not require the thinker to have any distinctive relationship to the object. Russellian definite descriptions, by contrast, do: they require us to think by means of a definite description ('the Queen of England', say) that picks out one object uniquely. Singular thought patterns with Russellian definites in requiring us to have a distinctive relationship to the object: for singular thought, it might be the sort of relationship that perception paradigmatically affords us. But unanchored discourse referents can be used to accumulate information about an object even by a thinker who has no distinctive relationship to the object. So their functional role is as *denoting concepts* that play a similar role to Fregean individual concepts. They allow us to think and talk about objects without being in a position to think singularly about them, or think about them by description.

Unanchored discourse referents can be introduced by introducing their truth-conditions. Let subscripted '*i*'s be terms for unanchored discourse referents.

(3a) A kitten was here. It was happy.

i_1 was a kitten. i_1 was here. i_1 was happy.

The discourse in (3a) is true at a maximally specific world-state w iff w contains some object that verifies the constraints on i_1 . It's true iff the world contains some object that was a kitten, was here, and was happy.

A belief ascription is true iff every *maximal extension* of the subject's beliefs *verifies* the constraints on the discourse referents, and

²⁵Irene Heim (1982) and Hans Kamp (1981) both influentially developed these ideas, around the same time. The distinction between anchored and unanchored discourse referents is due to Kamp (ms).

there is some maximal extension that does.

(5) Jane believes that a kitten was here. She also believes that it was happy.

- i_1 was a kitten
- i_1 was here
- i_1 was happy

To find the maximal extensions, start with the propositions that Jane believes. A maximal extension is the result of adding propositions to what Jane already believes until you get a complete, consistent description of the world.²⁶ A maximal extension *verifies* the items in the bullet-pointed list if the objectual existential closure of that list is true at that maximal extension.

Maybe one maximal extension of Jane's belief's includes the propositions: <Tabby was a kitten>, <Tabby was here>, and <Tabby was happy>. That maximal extension verifies the bullet-pointed list, because the objectual existential closure of that list is true at that maximal extension. Another maximal extension might include the corresponding propositions about Miley. That maximal extension also verifies the bullet-pointed list.²⁷

So what is it for Bill and Jane to agree that a kitten was here and to agree that it was happy? It is for every maximal extensions of *each person's individual* beliefs to verify the constraints on the discourse referent i_1 :

- i_1 was here
- i_1 was happy

So believing singular propositions about different individuals is itself sufficient for their agreement.

²⁶There are, of course, lots of interesting technical questions about what to do when the agent's beliefs are inconsistent, questions taken up in this enormous literature (e.g., in Kamp and Reyle (1993)).

²⁷Importantly, though, Jane need not be thinking about Tabby or Miley in order for these maximal extensions to be relevant. The maximal extensions are what are left open by her beliefs. In some cases, then, it's precisely because she's never thought about Tabby or Miley that they count as among the maximal extensions of her beliefs. If she's never thought about them, she's never formed any beliefs about them that would eliminate them as candidates for the happy kitten who was here.

If we posit unanchored discourse referents, ambitious hybrid theorists have a crisp answer to the §3 problem. Jeremy and David's agreement that abortion is permissible consists in:

(DYNAMIC) Jeremy and David agree that abortion is compatible with what is i_1 , while desiring to do what is i_1 .

And DYNAMIC is true because every maximal extension of their attitudes are extensions where the constraints on the discourse referents are true. Different extensions matter for Jeremy than for David. But that's not surprising – in (3), different extensions matter for Bill than for Jane.

4.4 Generalizing

I'm in the course of arguing that ambitious hybrid theories can solve the problem about agreement from §3. Up to now, I've been arguing by case, by considering different intensional anaphora. But it's possible to do even better. There is a simple and powerful argument that *any* adequate account of intensional anaphora is *guaranteed* to solve the problem about agreement.

The first step of the argument is to show that an account of intensional anaphora is adequate only if it vindicates the CENTRAL CONSTRAINT.

CENTRAL CONSTRAINT: Agents can share linked attitudes in virtue of having those attitudes to different particular individuals.

Suppose that Bill believes that Miley was a kitten who was here and happy, and Jane believes the same thing about Tabby. Then they agree that a kitten was here, and they agree that it was happy. This example establishes the CORE CONSTRAINT all by itself. It shows that individuals can have linked attitudes by having those attitudes to different particular individuals.

The second step is to show that any account that vindicates the CENTRAL CONSTRAINT will also solve the problems for the ambitious hybrid theorist. Remember that the problem was to explain how David and Jeremy could agree that abortion is permissible, where their agreement was grounded in attitudes about different

properties: *maximizing happiness* for Jeremy, *being-approved-by-an-impartial-spectator* for David. Ambitious hybrid theorists can and should gloss their agreement as linked attitudes: a belief and a desire, linked together, about the very same thing. The CENTRAL CONSTRAINT guarantees that there is some linked attitude that they will share, in virtue of having their particular attitudes.

Different accounts of intensional anaphora will offer different accounts of what the linked attitude is: whether it's an attitude about a singular proposition (as in §4.1), or an attitude about a descriptive proposition (as in §4.2), or an attitude about an unanchored discourse referent (as in §4.3). But the CENTRAL CONSTRAINT guarantees that there will be *some* linked attitude that they share in virtue of their distinctive belief/desire pairs. So it guarantees that there will be some solution to the problem from §3.

5 Does my solution undermine the interest of ambitious hybrid theories?

My general recommendation is that ambitious hybrid theorists defer to whatever account of intensional anaphora turns out to be correct. That is, they should formulate their views so that it can be implemented by drawing on the best explanation of intensional anaphora, whatever that is.

Now extant hybrid theorists have not conformed to this recommendation. So in making my recommendation, I am in effect charging that extant hybrid theorists have misunderstood which version of their view is best. Philosophers do sometimes fail to see which version of their view is best – it's plausible, for example, that Kripke (1975) failed to endorse the best version of his theory of truth.²⁸ But my charge is a still a very striking one to level, and a charge that calls out for very substantial defense.

To defend this charge, I need to do two things. First, I need to show that extant versions of ambitious hybrid theories face problems that my recommendation fixes. I think I've already done that: it should be common ground that ambitious hybrid theories struggle to explain agreement, and I think I've shown how my recommendation allows them to do better. Second, though, I need to show that hybrid

²⁸See pp. 170ff of Soames (1999) for discussion.

theorists can retain all the advantages that they care about if they adopt my recommendation.

Now there isn't space to consider each advantage that hybrid theorists have claimed. I'll rather give a general argument that they'll still retain those advantages after adopting my recommendation. Imagine a population where everyone has a range of moral judgments without speaking a language that allows them to express those judgments. For example, members of the community judge that killing is wrong, but they do not speak a language with any sentence that's synonymous with 'killing is wrong'. Their moral judgments still would be linked to motivation; they'll be motivated not to kill. And the moral judgments will still stand in inferential relations to each other. Hybrid theorists can offer the same powerful explanations as before, that the moral judgments consist of hybrids of beliefs and desires.²⁹

Now there would be several different ways for the population to learn to use language to communicate their moral judgments. One way is for their learning to use 'killing is wrong' to express the very belief/desire pair that grounds their judging that killing is wrong. Then David would use the sentence to communicate his belief that

²⁹There is an important complication here, because of the significant heterogeneity among hybrid theorists. Some of them take both mental states to be *semantically expressed* by attitude attributions. Others, by contrast, only take one of them to be semantically expressed, and take the other to be pragmatically or conventionally implicated.

Any of those hybrid theories should agree with the points in the previous paragraph. Those accounts are, among other things, accounts of what *attributors* take moral judgments to involve. Take Mary, a member of the community. Hybrid theorists are all giving accounts of what an utterance of 'Mary judges that killing is wrong' communicates: it communicates that she has both a belief and a desire. They differ in their accounts of how those states are communicated. But since they agree that both states are communicated, they will take Mary to have both of them if she has moral *target* beliefs. And they will agree on that point even in those cases where the attributee (Mary) wouldn't use the sentence.

Imagine someone who has an accident that eliminates their ability to produce language but leaves their moral judgments intact. Hybrid theories are still presumably accounts of that person's moral judgments. For example, they're still accounts of why that person is motivated to act in some ways and not others – because the person still judges that the action is morally forbidden, which is constituted by a belief and a desire. They offer different accounts of why we, as attributers, take the agent's judgment to be associated with the mental states – but those accounts should not depend on the attributee's ability to use moral language.

impartial spectators disapprove of killing, and the desire to do what an impartial spectator approves. Another set of ways would be for the population to learn to use ‘killing is wrong’ in one of the ways from §4. Maybe, for example, they learn to use it to communicate belief/desire pairs about unanchored discourse referents.

Ask yourself if learning to talk in one of the §4 ways would eliminate the basic mental states that constitute the population’s initial moral judgments. Suppose that it wouldn’t. Then all the §4 suggestions will preserve the philosophical advantages that ambitious hybrid theorists have cared about. Those mental states are themselves sufficient for the philosophical advantages that the hybrid theorists care about. So if the §4 suggestions wouldn’t eliminate them, the §4 suggestions wouldn’t change the philosophical advantages of the hybrid theory. And it *is* very hard to see why the choice of a way of talking would eliminate the population’s initial mental states. So it’s very hard to see why the §4 suggestion would undercut the philosophical advantages that the hybrid theorist claims.

In general, hybrid theorists should be open to different ways that moral discourse might implement their core idea. And they should see strong reasons for adopting one of the §4 ideas. §4 showed that natural language already has the resources to implement the core hybrid idea to crisply explain moral agreement. The simplest explanation of the facts about agreement is that each natural language implements hybrid views by using resources it already contains. Since hybrid theorists *can* adopt that simplest hypothesis, they should.

There is another important reason for ambitious hybrid theories to adopt one of the §4 suggestions. There is a plausible PUBLICITY constraint on what’s semantically expressed.

(PUBLICITY) Speaker-hearers know what S semantically expresses in virtue of their linguistic competence.

Now if the ambitious hybrid theorist is right, speaker-hearers can’t know which belief/desire pair Jeremy has in virtue of their linguistic competence. (It could be about ideal observers, or about happiness, as far as their linguistic competence goes.) Given the PUBLICITY constraint, propositions directly about his belief/ desire pair aren’t candidates for what’s semantically expressed. By contrast, the candidates from §4 do satisfy the PUBLICITY constraint. So PUBLICITY is another reason to favor them.

6 Further case studies

This paper has been narrowly focused on one particular challenge to hybrid theorists: a challenge about agreement. I've been focusing on the challenge about agreement because I think solving it is the key to solving a range of other pressing challenges.

This section suggests (without showing) that a range of other problems disappear if hybrid theorists adopt any of the suggestions from §4. I'll focus here on a dynamic account of intensional anaphora, with unanchored discourse referents. I'm focusing on them because they're the simplest option: they smoothly predict just the right truth-conditions for intensional anaphora without further complicated additions, like extensive use of situations.

Let's start by observing that ambitious hybrid theorists face problems about disagreement.³⁰ Imagine that Jeremy says that abortion is permissible. David hears what he said and disagrees, saying 'abortion is impermissible, so what Jeremy said is false!' To explain this discourse, we need to explain what the phrase 'what Jeremy said' refers to. A traditional and powerful suggestion is that 'what Jeremy said' refers to the representational proposition that Jeremy asserted. But it's hard for hybrid theorists to identify a representational proposition that can capture Jeremy and David's disagreement. For example, they needn't be disagreeing about the proposition that abortion is compatible with always acting to maximize happiness. David could grant that proposition, while still insisting that abortion is impermissible.

But 'said' is an attitude verb, like 'agree'. Ambitious hybrid theorists should offer a hybrid gloss. Jeremy's saying that abortion is permissible is a hybrid of saying something while also desiring something. Given the dynamic account of intensional anaphora, his saying that would be a hybrid of *saying* that abortion is compatible with what is i_1 while *desiring* to do what is i_1 .

So David's saying *abortion is impermissible* is a hybrid of saying that abortion is not compatible with what is i_1 while desiring to do what is i_1 . Given what he's said and what Jeremy has said, he's perfectly entitled to infer that what Jeremy said is false. If David is right that abortion is incompatible with what is i_1 , then a proposition that Jeremy said is false: the proposition that abortion is compatible

³⁰Schroeder (2009) is a *locus classicus* of the problem.

with what is i_1 .³¹

Now in offering this explanation, I need to explain why David and Jeremy's utterances are constrained to be about the same unanchored discourse referent. But we already have an explanation, from the end of §5. The only thing that competent speaker-hearers know *in virtue of their linguistic competence* about Jeremy and David is that they're both intending to talk about the demands of morality. So given the PUBLICITY constraint, the propositions they each express *have* to be about the same unanchored discourse referent.

Another important problem for ambitious hybrid theories is to explain inferences like the following.

1. Jefferson Davis judged that helping fugitive slaves was wrong.
2. But helping fugitive slaves wasn't wrong.
3. So Davis judged but didn't know that helping fugitive slaves was wrong.

This inference seems paradigmatically sound. But Brian Weather-son (2008) has argued that ambitious hybrid theorists can't explain why. An ambitious hybrid theorist associates the inference with the following belief/desire pairs:

- i.

BELIEF: Davis judged that helping fugitive slaves was K and desires to do what is K.
DESIRE: [none]
- ii.

BELIEF: Helping fugitive slaves wasn't J
DESIRE: to do what is J
- iii.

BELIEF: Davis judged but didn't know that helping fugitive slaves was K and desires to do what is K.
DESIRE: ???

³¹There are pressing foundational questions here, about propositions in a dynamic account. But if those foundational questions can't be answered, their not being answered just shows that some other account of intensional anaphora is correct. And I can use that other account to make the same claims I'm making here. Think, for example, of what proponents of liberalized singular thought would say about the propositions under discussion.

Crucially, the conclusion (iii.) has to be about the property that Davis is thinking about, the property K.³² But there isn't any valid inference to iii. from ii. and i. For one thing, Davis might be perfectly correct that helping fugitive slaves was K, and indeed know that it is, even if helping fugitive slaves isn't J.

Unanchored discourse referents have just the right structure to solve this problem. The ambitious hybrid theorist can associate the inference with linked thoughts about the very same unanchored discourse referent:

- a.

BELIEF: Davis judged that helping fugitive slaves was forbidden by what is i_1 and desires to do what is i_1 . DESIRE: [none]
--
- b.

BELIEF: Helping fugitive slaves wasn't forbidden what is i_1 DESIRE: to do what is i_1

- c.

BELIEF: Davis judged but didn't know that helping fugitive slaves was forbidden by what is i_1 and desires to do what is i_1 . DESIRE: to do what is i_1

My recommendation thus allows ambitious hybrid theorists to vindicate the soundness of this inference. More needs to be said, but again, the hybrid theorist's position is much improved.

Consider one last problem. Mark Schroeder (2009, 278ff) objects that ambitious hybrid expressivists violates the PUBLICITY constraint.

(PUBLICITY) Speaker-hearers know what S semantically expresses solely in virtue of their linguistic competence

And traditional formulations of ambitious hybrid theories do violate it, as noted in §5.³³

³²Weatherson emphasizes that context-sensitive terms make the same contribution in a knowledge report as they would in the corresponding belief report (Weatherson 2008, p. 539).

³³And he shows that ambitious hybrid theorists face similar problems given any standard account of the *expression* relation.

But unanchored discourse referents allow ambitious hybrid theories to conform to PUBLICITY. Both David and Jeremy could utter *abortion is permissible* to express approval of doing what is i_1 while also expressing that abortion is compatible with what is i_1 . So Jeremy and David's utterances are both associated with the very same things, and they can adopt any viable account of the *expression* relation. For example, they can adopt the Gricean (1957) idea that in making that utterance, Jeremy intends to cause his hearer to infer that he approves of doing what is i_1 and believes that abortion is compatible with what is i_1 .

7 Conclusion

There unfortunately isn't space here to dig into all the problems for ambitious hybrid theorists in the detail they deserve. I hope I've said enough to illustrate the fruitfulness of the ideas from §4, how accounts of intensional anaphora can also solve other problems besides the initial problem about agreement. It's legitimate for ambitious hybrid theorists to adopt any of those accounts, because they can adopt any of them without abandoning any of the view's interesting philosophical upshots. So I think I've given substantial evidence about the way forward for ambitious hybrid theorists: a user's guide for the tools that they are developing.

I've also given substantial evidence about what it takes to evaluate ambitious hybrid theories. You need to evaluate the version of the view that incorporates whatever account of intensional anaphora is correct. So I think that this paper has raised the bar for objections to ambitious hybrid theories. In order for an objection to be

A hybrid theorist might respond to Schroeder's criticism by arguing the competent hearers at least know that the speaker has beliefs that play a particular functional role, by being connected to desire. And the theorist might insist that the contents of the beliefs are not semantically expressed; the only thing semantically expressed is the functional role that the beliefs play. Teemu Toppinen (2013) and Mark Schroeder (2013) have developed views with this structure. As noted earlier, I do not think these views are helpfully classified as *ambitious* hybrid views. They do not allow that one speaker's use of a moral sentence can semantically express something different from what another speaker's use would express. Such views encounter a range of distinctive problems on their own – and the question in this paper is whether ambitious hybrid theories face severe enough problems that hybrid theorists should abandon them in favor of cautious views.

compelling, I think it needs to apply to hybrid accounts that adopt one of the suggestions in §4.

At at this point, I think we can be cautiously optimistic that even ambitious hybrid views can be made to work. So we should be ready to take them very seriously, both in metaethics and beyond.³⁴

References

- Alm, David. 2000. “Moral Conditionals, Non-Cognitivism, and Meaning.” *Southern Journal of Philosophy* 38:355–377
- Bar-On, Dorit and Chrisman, Matthew. 2009. “Ethical Neo-Expressivism.” In Russ Shafer-Landau (ed), *Oxford Studies in Metaethics, V. 4*, 133–166. (Oxford: Oxford University Press)
- Barker, Stephen. 2000. “Is Value Content a Component of Conventional Implicature?” *Analysis* 60:268–279
- Berker, Selim. 2017. “The Unity of Grounding.” *Mind*
- Björnsson, Gunnar and Finlay, Stephen. 2010. “Metaethical Contextualism Defended.” *Ethics* 121 (1):7–36
- Blackburn, Simon. 1998. *Ruling Passions*. (Oxford: Oxford University Press)
- Boisvert, Daniel. 2008. “Expressive-Assertivism.” *Pacific Philosophical Quarterly* 89:169–203
- Cappelen, Herman and Hawthorne, John. 2009. *Relativism and Monadic Truth*. (Oxford: Oxford University Press)
- Copp, David. 2001. “Realist-Expressivism: A Neglected Option for Moral Realism.” *Social Philosophy and Policy* 18:1–43
- Dreier, Jamie. 2005. “Meta-Ethics and the Problem of Creeping Minimalism.” *Philosophical Perspectives* 23–44
- Elbourne, Paul. 2005. *Situations and Individuals*. (The MIT Press: Cambridge, MA)

³⁴Special thanks to Steve Finlay, John Hawthorne, Barry Schein, Ralph Wedgwood, Jon Wright, to the referees and editors at *Mind*, and especially to Mark Schroeder and Scott Soames for many vigorous and illuminating discussions.

- Evans, Gareth. 1977. "Pronouns, Quantifiers, and Relative Clauses (I)." *Canadian Journal of Philosophy* 7:467–536
- . 1980. "Pronouns." *Linguistic Inquiry* 11:337–362
- Finlay, Stephen. 2014. *Confusion of Tongues: A Theory of Normative Language*. (Oxford: Oxford University Press)
- . 2016. "Disagreement Lost and Found." *Oxford Studies in Metaethics* 12
- Geach, Peter. 1967. "Intentional Identity." *Journal of Philosophy* 64(20):627–632
- Grice, H. P. 1957. "Meaning." *The Philosophical Review* 66 (3):377–388
- Hare, R. M. 1952. *The Language of Morals*. (Oxford: Oxford University Press)
- Hawthorne, John & Manley, David. 2012. *The Reference Book*. (Oxford: Oxford University Press)
- Heim, Irene. 1982. *The Semantics of Definite and Indefinite Noun Phrases*. Ph.D. thesis, University of Massachusetts
- Jackson, Frank. 1999. "Non-Cognitivism, Validity, and Conditionals." In Dale Jamieson (ed), *Singer and His Critics*, 18–37. (Oxford: Blackwell)
- Jeshion, Robin. 2010. "Singular Thought: Acquaintance, Semantic Instrumentalism, and Cognitivism." In Robin Jeshion (ed), *New Essays on Singular Thought*, 105–140. (Oxford: Oxford University Press)
- Joyce, Richard. 2001. *The Myth of Morality*. (Cambridge: Cambridge University Press)
- . 2006. *The evolution of morality*. (Cambridge, MA: MIT Press)
- Kamp, Hans. ????. "Proper Names from a communication-theoretic Perspective"
- . 1981. "A Theory of Truth and Semantic Representation." In Groenendijk Janssen and Stokhof (eds), *Formal Methods in the Study of Language*. Mathematical Centre, Amsterdam

- Kamp, Hans and Reyle, Uwe. 1993. *From Discourse to Logic Introduction to Modeltheoretic Semantics of Natural Language, Formal Logic and Discourse Representation Theory*. (Netherlands: Springer)
- Kaplan, David. 1989. “Demonstratives.” In Joseph Almog, John Perry, and Howard Wettstein (eds), *Themes from Kaplan*. (New York: Oxford University Press)
- Kripke, Saul. 1975. “Outline of a Theory of Truth.” *Journal of Philosophy* 72:690–716
- Lewis, David. 1975. “Adverbs of Quantification.” In Edward L. Keenan (ed), *Formal Semantics of Natural Language*. (Cambridge: Cambridge University Press)
- MacFarlane, John. 2014. *Assessment Sensitivity: Relative Truth and its Applications*. (Oxford: Oxford University Press)
- Moore, G E. 1922. “Philosophical Studies.” In *Philosophical Studies*. London: Routledge & Kegan Paul.
- Moss, Sarah. 2017. *Probabilistic Knowledge*. (Oxford: Oxford University Press)
- Perl, Caleb. 2017. *Positivist Realism*. Ph.D. thesis, University of Southern California.
- . forthcoming. “Might moral epistemologists be asking the wrong questions?” *Philosophy and Phenomenological Research*
- Ridge, Michael. 2006. “Ecumenical Expressivism: Finessing Frege.” *Ethics* 116:36–302
- . 2014. *Impassioned Belief*. (Oxford: Oxford University Press)
- Ross, W D. 1930. *The Right and the Good*. (Oxford: Oxford University Press)
- Santorio, Paolo. 2016. “Nonfactual Know-How and the Boundaries of Semantics.” *Philosophical Review* 125(1):35–82
- Schroeder, Mark. 2009. “Hybrid Expressivism: Virtues and Vices.” *Ethics* 119 (2):257–309

- . 2011. “\textit{Oughts}, Agents, and Actions.” *Philosophical Review* 120 (1):1–41
- . 2013. “Tempered Expressivism.” In *Oxford Studies in Metaethics* 8, 283–314. (Oxford: Oxford University Press)
- Sidgwick, Henry. 1907. *The Methods of Ethics*. (Indianapolis, IN: Hackett)
- Silk, Alex. 2016. *Discourse Contextualism: A Framework for Contextualist Semantics and Pragmatics*. (New York: Oxford University Press)
- Smith, Michael. 1994. *The Moral Problem*. (Oxford: Blackwell)
- Soames, Scott. 1999. *Understanding Truth*. (Oxford: Oxford University Press)
- Stalnaker, Robert. 1998. “On the Representation of Context.” *Journal of Logic, Language, and Information* 7:3–19
- Stevenson, Charles Leslie. 1937. “The Emotive Meaning of Ethical Terms.” *Mind* 46 (181):14–31
- Strandberg, Caj. 2012. “A Dual Aspect Account of Moral Language.” *Philosophy and Phenomenological Research* 84 (1):87–122
- . 2015. “Options for Hybrid Expressivism.” *Ethical Theory and Moral Practice* 18 (1):91–111
- Toppinen, Teemu. 2013. “Believing in Expressivism.” In Russ Shafer-Landau (ed), *Oxford Studies in Metaethics* 8, 252–282. (Oxford: Oxford University Press)
- Weatherson, Brian. 2008. “Attitudes and Relativism.” *Philosophical Perspectives* 22:527–544
- Yalcin, Seth. 2007. “Epistemic Modals.” *Mind* 116:983–1026