

The Pragmatics of Normative Disagreement

Stephen Finlay

This is a preprint version of an article to appear in G. Fletcher & M. Ridge (eds.), *Janus-Faced Judgment: Hybrid Approaches to Meta-Normative Theory*. New York: Oxford University Press, 2014. It is based on Chapter 8 of my book, *Confusion of Tongues: A Theory of Normative Language* (New York: Oxford University Press, 2014), and is revised for length and to stand alone.

On their surface, moral and normative sentences appear descriptive. But what kinds of facts could they describe? Many have been drawn to the idea that normative words like ‘ought’ and ‘good’ are used to refer to *relational* properties, consisting in a relation to something like a standard, rule, desire, or end, which can vary between contexts. On my preferred view, to say that *s* ought to φ is roughly to assert the proposition that some implicit *end, e*, is more likely to obtain if *s* φ s than if *s* does anything else instead.

But relativistic metaethical theories face a familiar, seemingly fatal problem in accounting for the intuitive extent of normative disagreement (and agreement). They imply that superficially incompatible moral or normative claims made by speakers who are concerned with different ends (etc.) express logically consistent propositions. Given the natural idea that disagreement involves logical inconsistency, these theories seem to imply that utterances of sentences like ‘*s* ought to φ ’ and ‘*s* ought not to φ ’ don’t really express disagreements even when they intuitively do.

In particular, moral disagreements evidently occur between people who aren’t united by preference for any common end; call such disagreements *fundamental*. It’s widely thought that accommodating these disagreements requires instead some kind of *absolutist* semantics: either a form of nonrelational descriptivism identifying a common property of concern to disagreeing parties, or a form of expressivism identifying some kind of disagreement in attitudes other than belief, like preferences.

I’ll first explore a different source of disagreement problems. There are compelling reasons for thinking that normative propositions are information-relative, which generates structurally identical problems about disagreement and shows that normative disagreement is everybody’s problem. I’ll argue that any adequate account must appeal to pragmatics, and that the end-relational theory (and so any relevantly similar view) provides a “quasi-expressivist” solution in terms of the conflicting preferences that are pragmatically expressed by asserting logically consistent normative propositions—a solution that extends straightforwardly from information- to end-relativity.

So far, this recapitulates my moves in a previous paper, cowritten with Gunnar Björnsson (2010). But whereas we there offered a pragmatic solution only for a specific kind of case involving a deliberating agent and a better-informed advisor, here I’ll advance a general theory of normative disagreement. I’ll also argue that despite being a purely descriptivist semantics, the end-relational theory actually provides a better account of disagreement in attitude than expressivism, making it redundant to posit expressive conventions in the meanings of normative words. Contrary to the

received wisdom, intuitions of normative disagreement support a relational semantics over either expressivism or nonrelational descriptivism. For reasons of space, however, I won't be able to discuss related issues about how normative claims are reported and evaluated as "true" or "false".

1. The Problem of Disagreement from Information-Relativity

An initial, naïve formulation of the natural idea that disagreement requires inconsistent propositions is

Inconsistent Belief: B's assertion/belief that p disagrees (agrees) with A's assertion/belief that q iff p and q are logically inconsistent (equivalent).

If an assertion or belief of B's disagrees with one of A's, then by extension B himself can be said to disagree with A herself in that respect. This can be said in a familiar sense even where there isn't any interaction between them, though talk about "disagreement" may sometimes suggest it.

Consider the following scenario¹:

Two Gun Roulette. Angie has been kidnapped by Sadie, who tells her that Angie will be tortured to death unless she plays a special game of Russian Roulette. She must choose between two six-chambered revolvers, L to her left, and R to her right. L has been loaded randomly with one round, R with three rounds. Angie must aim her chosen revolver at her head and pull the trigger once. If she survives she will be released unharmed. She has no other realistic options.

Evidently, Angie could appropriately conclude her deliberations by saying,

A: 'I ought to use L, not R.' (*Deliberation*)

The end-relational semantics (ERT) accommodates this as assertion of the true proposition that she is more likely relative to her information to survive unharmed if she uses L than if she uses R. Now let's supplement the scenario:

Advice: Bertie is a fellow captive who has all Angie's information plus the following: since the revolvers were loaded, five captives have played Sadie's game. (In this version, the cylinders are spun once upon loading but not between each use.) All five chose L and survived. Bertie has one opportunity to whisper briefly in her ear, but not enough time to share this information.

The implication of Bertie's information is that the next use of L will be lethal. Evidently he would speak appropriately by saying,

B: 'You ought to use R, not L.' (*Advice*)

¹ This is a variant of the miner case in Regan 1980, Kolodny & MacFarlane ms. I introduce *Two Gun Roulette* to address the problem relative to a single end, which avoids complications for ERT addressed in Finlay 2014, Ch. 6.

ERT accommodates this as assertion of the true proposition that Angie is more likely relative to Bertie's information to survive unharmed if she uses R than if she uses L.

To accommodate these claims, ERT apparently must interpret them as relativized to different information. Consequently, Angie and Bertie assert logically consistent propositions. It further seems they needn't have any inconsistent beliefs. Bertie accepts the proposition assigned as the content of Angie's claim, and Angie doesn't believe anything inconsistent with the content assigned to Bertie's claim. By *Inconsistent Belief* they therefore don't disagree, yet intuitively Bertie does disagree with Angie's statement. (It isn't so intuitive that Angie disagrees with Bertie; we'll return to this asymmetry below). He seems able to say appropriately, 'No, you ought to use R, not L,' and while ERT suggests that Angie asserted a proposition that Bertie may know to be true, for him to say 'Yes, that's right. You ought to use L, not R' would clearly be perverse and deceptive. Here we confront a puzzle: Angie and Bertie seem to assert propositions whose truth-conditions are determined wholly by their own differing information, and are therefore logically consistent, but Bertie seems thereby to disagree with Angie.

This puzzle arises for any theory that relativizes normative propositions to different information. One response simply denies that these normative claims really are information-relative, or interprets both as relativized to something like *all the facts*.² On this view, Angie asserts a proposition concerning what she *objectively* ought to do, which is simply false though she may be justified in believing it, and Bertie is therefore correct to reject it. But this analysis can't be correct, because Angie may know it's quite likely she objectively ought to use R. She can appropriately say 'I ought to use L, although it's quite likely that given all the facts I ought to use R,' but she can't say 'Given all the facts I ought to use L, although it's quite likely that given all the facts I ought to use R,' since accepting that *p* is likely false is incompatible with warranted assertion of it. Angie thereby can't be using 'ought' in an objective sense.³ Similar reasoning shows that Bertie also can't be using 'ought' objectively: he knows that using R is quite likely also to be fatal, which would also make it false that Angie objectively ought to use R. Making sense of these claims requires recognizing their relativity to the speakers' incomplete information.

A second response is that Angie's and Bertie's statements are relative to the same incomplete information. Deliberating agents presumably don't aim to make decisions on the basis of merely the information in their possession when they begin deliberating, but on roughly *the fullest information they can utilize at the time of decision*, which will include information others are in a position to make available to them. Whereas Bertie is unable to share his information with Angie, information can be "available" for use without being possessed, if others communicate the probability of *p* relative to it. Angie's use of 'ought' can therefore be expected to be relativized to information selected this way

² Judy Thomson holds that all 'ought' claims are sensitive to ("objective") probability relative to the *information available to humans at the time* (2008: 195). *Two Gun Roulette* is inspired by her scenario.

³ This is clearer where the speaker knows one of the other options is objectively best, like the miner case; see Kolodny & MacFarlane ms.

rather than merely to her own information: call this *news-sensitivity*.⁴ Her statement would then be false because sensitive to Bertie’s additional information, although perhaps still justified since she lacks reason to suspect this information exists. Bertie’s statement would directly contradict hers, accounting for disagreement by *Inconsistent Belief*. But while news-sensitivity seems plausible for cases of advice like Bertie’s, it isn’t a sufficiently general solution. If Bertie were rather a long-distance eavesdropper, intuitively he could still disagree with Angie though his information couldn’t be available to her.⁵ To simplify discussion I’ll proceed under the assumption that news-sensitivity doesn’t account for disagreement in the original case either.

A third response rejects these ordinary intuitions as mistaken: Bertie doesn’t genuinely disagree with Angie, because their claims are consistent. The intuitions result from failing to appreciate that the claims are implicitly relativized to different information, as Bertie does himself. If intuitions of normative disagreement are systematically unreliable then the entire problem could be dismissed as based on faulty data. But this proposal fails to account for the subtlety of ordinary intuitions of normative disagreement. Not just any two speakers asserting superficially incompatible normative sentences are taken to be disagreeing—such claims are often recognized as talking past each other—and we’ll see that these intuitions exhibit a high level of sensitivity to information-relativity. (I take no stand on what should be classified as “genuine” disagreement; the task is rather to explain the *intuitive sense* of disagreement.)

These cases of disagreement from diverging backgrounds thus pose a general puzzle for metaethics—or rather for philosophy in general, since analogous problems are observed for many words that encourage relativistic treatments, such as language about possibility (‘might’), knowledge (‘knows that’), taste (‘disgusting’), and gradable adjectives (‘tall’). E.g.:

Angie: ‘I might survive if I use L.’

Bertie: ‘No, you cannot survive if you use L. But you might survive if you use R.’

Sadie: ‘No, you’re both wrong. Angie cannot survive no matter which she uses.’⁶

A fourth response to these problems is *relativism about truth*.⁷ The idea is that claims made relative to some k don’t assert propositions about relations to k , but absolutist propositions that only have truth-values relative to perspectives characterized by indices of k ’s kind. So normative claims aren’t information-relative in content, but instead are true or false only relative to particular information-perspectives. Angie and Bertie would then both address the same simple proposition that *Angie ought to use L and not R*, which isn’t true or false simpliciter but only relative to different perspectives. Angie correctly asserts it, as true relative to her information. Bertie correctly rejects it, as false relative to his. Because Bertie rejects the same proposition Angie asserts, truth-relativism secures the result that they satisfy *Inconsistent Belief*. But this solution isn’t as straightforward as advertised.

⁴ Dowell 2013 argues this is a sufficient solution.

⁵ See also other cases below, particularly that Connie can appropriately either agree (postmortem) or disagree (hindsight) with Angie.

⁶ Suppose Sadie knows the next use of either will be lethal.

⁷ E.g. Kölbel 2002, MacFarlane 2007.

Truth-relativism needs to appeal to pragmatic resources which are also sufficient for a solution on our contextualist semantics.

2. Disagreement as Pragmatic

Truth-relativism's claim to account for the sense that Bertie disagrees with Angie is based on satisfying *Inconsistent Belief*. But this is only sufficient for disagreement given a traditional, nonrelativist view of truth. Consider

Postmortem: Connie is a detective investigating after the fact. As it happened Angie didn't hear Bertie's whisper and played using L, shooting herself in the head. By an incredible fluke, the bullet passed through her skull nonfatally, leaving Angie with a severe injury. Connie has interviewed Bertie and learned all his information. She visits Angie in hospital weeks later.

On truth-relativism, the proposition Angie asserted by saying 'I ought to use L, not R' is false relative to Connie's information, which is relevantly similar to Bertie's. Connie knows this, but can agree with Angie's claim, unlike Bertie, with a (figuratively) *postmortem* claim:

C: 'You were right. You ought to have used L, not R. In light of what you knew it was the only sensible action.' (*Postmortem*)

This case behaves as we'd naïvely expect if contextualism rather than truth-relativism were correct. As contextualism also predicts, Connie wouldn't seem to be disagreeing with Bertie. She could alternatively say, 'Bertie was right, you ought to have used R'.

Consider also

Less Informed: Debbie is Angie's friend, captured with her. Debbie has less information than Angie, not knowing how many rounds were loaded into L or R, though she knows Angie knows.

On truth-relativism, Angie's claim that she ought to use L is false relative to Debbie's inferior information. Debbie knows this, but doesn't intuitively disagree with Angie's claim.⁸ It would obviously be inappropriate for her to say, 'No, you're wrong. Neither option is any better than the other.' She can rather say,

D: 'I don't know whether you ought to use L or R.' (*Less Informed*)

This identifies a general problem for truth-relativism. In its least ambitious incarnations, truth-relativism holds that truth is relative to a *time* of assessment, which it therefore excludes from the propositions. But consider:

⁸ Perhaps Dora can infer from Angie's claim that Angie is more likely to survive if she uses L, which makes Angie's claim true relative to Dora's information too. But suppose Dora also hears Bertie speak without knowing which is better informed: still she intuitively wouldn't disagree with Angie.

Déjà Vu: One week after surviving Sadie's game, Angie has the terrible luck to be kidnapped by Sadie again, who compels her to play the game with the same revolvers, but reversed so R holds the single round. Eddie is another captive who knows the next use of R will be lethal, so he asserts, 'You ought to use L, not R.'

According to this basic version of truth-relativism, Eddie accepts as true (relative to his own temporal perspective) the same proposition Angie asserted, and rejects as false the proposition Bertie asserted. He thereby satisfies the *Inconsistent Belief* criterion for agreeing with Angie's claim and disagreeing with Bertie's. But intuitively he doesn't seem to be talking about the same thing, as a contextualist treatment of time-relativity would predict.

These cases illustrate that asserting superficially incompatible sentences, like 's ought to φ ' and 's ought not to φ ', isn't sufficient to trigger intuitions of disagreement. Contextualist theories provide a simple and conservative explanation: the same normative sentence is used to express different propositions. But truth-relativists are obliged to conclude that satisfying *Inconsistent Belief* isn't sufficient for intuitively disagreeing, and therefore concede the story to be more complex and to involve further conditions. While contextualist theories are still in trouble provided inconsistent beliefs are necessary for disagreement, I'll argue that the missing conditions are sufficient by themselves and without inconsistent belief.

These intuitions of disagreement are evidently responding to differences between the contexts of *Advice*, *Postmortem*, and *Less Informed*, rather than simply the sentences used. Two kinds of variability can be ruled out. First, while the speaker's *information* may vary between contexts, Bertie and Connie possess relevantly identical information yet Connie but not Bertie can appropriately agree with Angie. Second, while the speaker's *temporal relation* to the original claim may also vary (Bertie speaks prospectively, Connie retrospectively), notice that Sadie can felicitously agree prospectively with Angie, though occupying the same temporal and informational perspective as Bertie:

S: 'Yes, you're right. You ought to use L, not R.' (*Pseudo-Advice*)

Since Sadie doesn't want Angie to survive unharmed she isn't interested in making her fuller information *i* available for Angie by telling her what she ought given *i* to do. But she can still sincerely express agreement with Angie or Connie about what ought to be done relative to Angie's inferior information.

The relevant factor is evidently the *conversational ends* of the speakers, as truth-relativists themselves acknowledge: Mark Richard (2011) argues that whether an assertion that *p* disagrees with an assertion that not-*p* depends on its point, and Niko Kolodny and John MacFarlane (ms) maintain that a speaker like Bertie must be disagreeing with the original claim because he is offering *advice* which by nature aims to help the agent solve the practical problem she deliberates over. In other words, the sense of disagreement here depends on recognizing that Bertie speaks *cooperatively*, making it sensitive to pragmatic and not merely semantic cues.

This role of pragmatics in generating disagreement intuitions can also be demonstrated by observing a flaw in our naïve formulation of *Inconsistent Belief*. Speakers can evidently disagree without asserting contradictory propositions. Suppose A says, ‘We ought to go stargazing at 8pm’, and B replies, ‘Daylight savings started today.’ B is naturally interpreted as disagreeing with A, though the propositions asserted aren’t inconsistent. The disagreement lies not in what is said but in what is pragmatically expressed. Roughly, in saying that p , B disagrees with A by pragmatically expressing the belief that *not-q*, by virtue of what p entails in combination with the common ground. In this example, p (daylight savings started today) implies *not-q* (it’s not true they ought to go stargazing at 8pm) on the assumptions that they ought to go stargazing only after dark and that the sky isn’t dark at 8pm during daylight savings. This suggests a revised principle:

*Inconsistent Belief**: B’s assertion that p disagrees (agrees) with A’s assertion that q iff B thereby expresses the belief that *not-q* (q).

The objection is that relational semantics can’t accommodate normative disagreements between advisors and agents, because the cooperative nature of advice means that it must be intended to help the agent reach the right answer to the question she asks in her deliberations about what she ought to do. This can’t be a question about what the agent ought to do relative to her own information, for which she neither needs nor receives any help, but it also can’t be about what she ought to do relative to the advisor’s information, since this wasn’t her concern in deliberating. I’ll argue that ERT answers this objection when we apply a single basic principle of pragmatics: that *we expect people to speak in the way they consider best for their conversational ends*.⁹ Where conversational ends are shared this implies a version of Grice’s cooperative principle, but we’ll also be interested in other kinds of context.

3. Information-Relativity in Contexts of Direct Interest

Begin with the pragmatics of agents’ claims, like Angie’s. This is a context of deliberation, so we must identify the motivations of deliberating agents. The objection assumes this to be the end of discovering which proposition of the form *I ought to φ* is true, but ERT suggests a simpler, more natural answer: agents deliberate with the aim of achieving their particular motivating ends. Angie’s ultimate goal in deliberating is simply *that she survives unharmed*.

Knowing the truth of an ‘ought’-proposition relativized to an end is of special instrumental interest to an agent in what I’ll call a context of *direct interest*, where the agent is actively motivated toward the end. Knowledge of what is most likely to achieve an end is an ideal basis for making a decision aimed at achieving that end. But this depends on the quality of the information on which this probability is based. An ‘ought’-proposition relativized to the fullest available information will provide the instrumentally best available basis for a decision; i.e. one that most increases the probability of the end relative to the available information. ERT therefore implies that an agent has

⁹ See Finlay 2014, Chapter 5 for discussion.

no better means from her own point of view for promoting her end through deliberation than by identifying what she ought to do relative to the fullest information available. So Angie would conclude her deliberation by judging what she ought to do relative to the information available to her, in order to survive unharmed.

Now consider the pragmatics of Bertie's utterance. This is a context of advice, where at least here¹⁰ the speaker is expected to cooperatively borrow the agent's motivating end with the aim of guiding the agent toward it. Bertie also therefore speaks in a context of direct interest in Angie's surviving unharmed. Because he is better informed, he knows what she ought to do in order to survive unharmed relative to fuller information. This truth is of greater instrumental value relative to their shared end than the truth she asserts, and so Bertie best promotes their shared conversational end by communicating the proposition relativized to his information.

Since Angie's fundamental interest is in surviving unharmed, and not in knowing the truth of the proposition she asserted, this is the most cooperative thing Bertie can do. He thereby makes his information available for her to act on; because it provides a better basis for a decision aimed at Angie's surviving unharmed,¹¹ she would prefer to act on the basis of this proposition rather than the one she asserted. So ERT accommodates the intuition that Bertie's advice is cooperative, even though he addresses a different proposition. It predicts that (i) this is the proposition he would assert; (ii) he wouldn't waste his or Angie's time by addressing the proposition she asserted, which his information makes moot, and also (iii) this is what Angie would expect him to do.

It's also now easy to explain why Connie, Debbie, and Sadie aren't similarly expected to deny that Angie ought to use L, although it's false that she ought to use L relative to their respective information. Unlike Bertie they wouldn't thereby promote their conversational ends, though for different reasons. While less-informed Debbie has direct interest in Angie's end of surviving unharmed, denying the proposition relativized to her lesser information wouldn't promote it in any way. On the contrary, because of the expectation of cooperation Debbie would thereby pragmatically indicate she has more information than Angie, which is likely to cause Angie to make a subjectively worse decision. By contrast, Connie and Sadie each speak in contexts of *indirect interest*, addressing a proposition relativized to an end toward which they're not actively motivated. Connie speaks in a postmortem context where achieving Angie's end is no longer at issue, while Sadie speaks in a noncooperative context of pseudo-advice, hoping that Angie's end doesn't obtain. They therefore aren't pragmatically obliged to address the proposition relative to the fullest available information, unlike Bertie and Debbie.

This explains the appropriateness of the 'ought' claims in the various scenarios, but doesn't yet explain intuitions of disagreement (e.g. that Bertie disagrees with Angie). But it provides the resources for a solution. In contexts of direct interest in *e*, asserting that *in order that e it ought to be that*

¹⁰ Advice sometimes guides agents toward ends they *would* prefer if better informed.

¹¹ The proposition that *relative to fuller information $i_{(+)}$ e is most likely given p* itself makes it true that *e* is most likely given *p* relative to the lesser information *i₍₋₎*. A subjectively better basis for decision doesn't necessarily result in an objectively better outcome, of course.

p (relative to the fullest available information) will pragmatically express a preference that p .¹² Angie’s utterance therefore expresses her preference that she uses L, not R, while Bertie’s utterance expresses his preference that Angie uses R, not L. By asserting consistent propositions they pragmatically express a common kind of attitude toward inconsistent propositions.

Could an expressed conflict of preferences be sufficient to trigger intuitions of normative disagreement? Expressivists since Charles Stevenson have labeled this kind of conflict a “disagreement in attitude”, in contrast to a “disagreement in belief”, and have used it to explain normative disagreement. It promises to accommodate the natural idea that disagreement involves inconsistency in attitudes, but without satisfying *Inconsistent Belief*, which motivates the objection. Whereas purely descriptivist semantics are generally assumed to be committed to explaining disagreement through some version of *Inconsistent Belief*,¹³ ERT suggests rather the quasi-expressivist solution that normative disagreement can occur through pragmatic expression of inconsistent attitudes of a different kind.¹⁴

4. Information-Relativity in Contexts of Indirect Interest

So far we only have an account of normative disagreement in one kind of context. A complete solution requires a general theory accommodating disagreement intuitions across a full range of contexts. If some normative disagreements consist in pragmatically expressed disagreement in preference then the simplest remaining theory is that all are; i.e.

Inconsistent Preference: B’s assertion that p normatively disagrees (agrees) with A’s assertion that q iff B thereby expresses a preference inconsistent with (equivalent to) the preference A thereby expresses.

As ERT maintains that normative utterances express beliefs in addition to any preferences, it may seem strange if no normative disagreements consisted solely in inconsistent beliefs; expressivists may complain that the descriptive semantics is a boondoggle doing no significant work.

Another possibility is that disagreement intuitions are sensitive to inconsistency either in belief or preference—which threatens to be messy, and might seem unnecessarily complex compared to expressivism. To avoid objections of being ad hoc, a mixed account must take the form of a general principle with testable predictions. While the theoretical space here is large, we won’t need to look far beyond the simplest, disjunctive version:

Inconsistent Belief or Preference: B’s assertion that p normatively disagrees (agrees) with A’s assertion that q iff B thereby expresses *either* (i) an end-relational belief inconsistent with (equivalent to) an end-relational belief A

¹² For discussion see Finlay 2014, Chapter 5.

¹³ E.g. Horgan & Timmons 1992.

¹⁴ See also Robinson 2009, who calls this “quasi-disagreement”.

thereby expresses, *or* (ii) a preference inconsistent with (equivalent to) the preference A thereby expresses.

On first glance, contexts of indirect interest favor a mixed account. Our quasi-expressivist treatment of the disagreement in preference between Angie (*Deliberation*) and Bertie (*Advice*) appealed to two contingent features of contexts of direct interest: the speaker's aiming at the salient end, and relativizing to the fullest information available. Normative speech and thought in which either of these features are absent, like Connie's (*Postmortem*) and Sadie's (*Pseudo-Advice*), won't in the same way express preferences on ERT, or thereby disagreement in preference. But Connie and Sadie both seem to agree with Angie when they say, 'You ought to use/ have used L, not R', which may seem to rule out *Inconsistent Preference*.

These results are independently plausible: in her postmortem context, Connie presumably prefers rather that Angie *didn't* use L, since she'd prefer Angie had been released unharmed. Because Sadie shares Bertie's information she may in fact prefer that Angie uses L, but her utterance won't express this since her audience has no reason to interpret it as relativized to her full information, given that she is offering uncooperative pseudo-advice. Their expressed agreement with Angie therefore couldn't consist in indicating that they share Angie's preference for her using L and not R.

By contrast, *Inconsistent Belief or Preference* accommodates these as agreements in belief. Connie and Sadie both know that Angie asserted a true proposition relativized to her own information, and we can straightforwardly explain the pragmatic appropriateness of agreeing with it in their contexts of indirect interest. Since Angie's surviving unharmed isn't their conversational end, they aren't pragmatically expected to address a proposition relativized to the fullest information available.

These speakers are interested instead in what deliberative conclusion Angie should reach in her epistemic situation. Connie's end might be to comfort Angie, by reassuring her that she couldn't reasonably be expected to have done anything differently. Sadie might find perverse amusement in "helping" Angie make a good decision within her epistemic limitations. In these contexts evaluating the information-relative proposition Angie asserted is directly relevant to the speakers' goals, so Connie and Sadie are each naturally interpreted as expressing agreement in belief.

Other disagreements in contexts of indirect interest are problematic for *Inconsistent Belief or Preference* as well as for *Inconsistent Preference*, however. So far we've found speakers in contexts of direct interest disagreeing in preference, and speakers in contexts of indirect interest disagreeing in belief. But sometimes a speaker in a context of indirect interest intuitively disagrees with a normative claim when our semantics predicts that no inconsistency in belief is expressed. One illustration is what I'll call a *hindsight* context, where a speaker addresses retrospectively what ought to have been done, in light of her present information.

Consider Connie again, who can not only appropriately agree postmortem with Angie's earlier claim, but in the very same circumstances could appropriately say instead,

C: 'It turns out you were mistaken. You ought to have used R, not L.' (*Hindsight*)

(Indeed, Angie herself may appropriately say at this later time, ‘In hindsight I was mistaken. As I now know, I ought to have used R and not L.’ Intuitively, she thereby expresses disagreement with her earlier deliberative statement.) But according to our theory, Connie knows that the belief Angie earlier expressed (that *she ought relative to her information to use L rather than R*) was true, so these hindsight claims don’t express disagreement in belief. Given that the theory also doesn’t predict that these claims express any preference, these disagreement intuitions are unaccommodated.

The problem can also be observed for agreement, from a variation on the hindsight context:

Hindsight Agreement: Freddie is a detective who investigates the revolvers after Angie survived her injury from using L, discovering that use of R would have been fatal. Freddie is misinformed on two counts: he was told that L was loaded with a blank so that Angie survived her ordeal unharmed, and also that Sadie lied to Angie that L held the three rounds and R only held one. He therefore believes that Angie was badly confused about her own information when she said she ought to use L.

In this context, Freddie can appropriately say in hindsight,

F: ‘Angie was right; she ought to have used L, not R. Though not for the reason she thought!’ (*Hindsight Agreement*)

Freddie seems to have agreed with Angie’s statement. But by ERT he must believe that the proposition she asserted is false, so this can’t be agreement in belief. If we predict that he also fails to express any preference then this agreement intuition is also unaccommodated.

The assumption that normative speech doesn’t express preferences in contexts of indirect interest should be reconsidered. Preferences might be expressed in some such contexts even if our previous explanation doesn’t apply. Indeed, the hindsight claims made by Connie and Freddie do plausibly express preferences. Regardless of their conversational ends, both can be presumed to prefer that Angie survived unharmed. Since these hindsight claims are relativized to the fullest information available to the speaker, they concern what was optimal from the speaker’s present point of view for achieving an end she prefers.

This solution isn’t sufficiently general, however, because preference for the salient end is contingent; for example, Sadie can also appropriately disagree in hindsight with Angie. As she is dragged off in handcuffs she might say, ‘You were wrong. You ought to have used R’, perhaps gloating over how she manipulated Angie into harming herself. On our theory this can’t be disagreement in belief as for the previous cases of hindsight, but since Sadie transparently prefers that Angie didn’t survive unharmed, her statement also doesn’t indicate she prefers that Angie used R; if anything it indicates the contrary preference.

This case of alienated hindsight therefore still resists both *Inconsistent Preference* and the disjunctive principle. Additionally, Connie’s and Sadie’s hindsight claims seem to express the same kind of disagreement with Angie’s deliberative claim, which suggests this isn’t the right solution also for the cases where the speaker is known to prefer the salient end. In general terms, pragmatic expression of preferences seems to depend on end-relational claims being relativized both to an end

the speaker prefers and to the fullest information available to the speaker, but intuitively normative disagreements arise between speakers with different preferences and information.

5. Attitudes with Conditional Content

If some intuitive cases of normative disagreement can't be explained as disagreements in either belief or attitude this poses a problem for any theory. The prospects for a solution appealing entirely to inconsistent beliefs look hopeless. But expressivists offer a way to account for these problematic cases as disagreements in attitude. Allan Gibbard proposes that the relevant attitudes aren't directed toward options simpliciter (e.g. that Angie uses R and not L) but toward *conditionals*, or options under particular conditions.¹⁵ A normative utterance of '*s* ought to φ ' expresses the speaker's attitude toward φ -ing *on the condition of being in *s*'s situation, or of being *s* herself*, i.e. toward the proposition expressed by 'I φ if I am *s*'. Sadie's disagreement with Angie might then be accounted for as follows: Sadie's hindsight statement that Angie ought to have used R and not L expresses her attitude toward using R and not L if she herself were Angie/ in Angie's situation *C*. Angie's deliberative statement that she ought to use L and not R expresses the same attitude toward using L and not R if in *C*.

This approach derives some intuitive appeal from its similarity to the following kind of exchange, which is naturally interpreted as a case of disagreement:

A: 'If *I* were *s*, I'd φ .'

B: 'Well if *I* were *s*, I'd ψ .'

It might also seem easily extended to disagreements or agreements in attitude expressed in the contexts *Advice*, *Postmortem*, and *Hindsight*, which would support some version of the simpler principle *Inconsistent Preference* and suggest that expressivism offers a simpler, better account of normative disagreement than any descriptivist semantics.

This strategy requires a different kind of attitude, as *preferences* toward conditionals generate the wrong results. Connie's postmortem claim that Angie ought to have used L can't express the preference that if she were in Angie's situation she uses L, since she knows that using L is worst relative to her preferred end of surviving unharmed.¹⁶ Gibbard therefore appeals instead to attitudes of *planning* or intention, which avoids this problem. Connie's utterance more plausibly expresses her plan to use L and not R if in Angie's situation. Since plans are designed to be implementable in the salient situation, Connie can only make plans for what to do if in Angie's situation that would make sense to her if she was indeed in Angie's situation—including being limited to Angie's lesser information. I'll temporarily switch to talking about plans, though we'll see this introduces new problems.

¹⁵ Gibbard 2003: 69.

¹⁶ It doesn't help to say that the relevant preference is that if she were in a situation *subjectively like* Angie's then she uses L rather than R, since Angie is in such a situation and Connie prefers that she uses R in it.

Another question concerns how to understand the idea of the speaker's "*being s*". Taken literally, worries arise about the coherence of this. Isn't Sadie's being Angie a metaphysical impossibility, and if she were Sadie wouldn't she do exactly the same thing as Sadie? These problems can be avoided by understanding the strategy in terms of *being in s's situation*: there is nothing absurd about Sadie's contemplating herself being "in Angie's shoes". But this move problematizes the claim of inconsistency in speakers' attitudes.

One perhaps minor problem is that different speakers' uses of 'I' refer to different people. The propositional content of Sadie's thought, *I use R and not L, if I am in C*, would therefore seem to be that *Sadie* uses L and not R if *Sadie* is in *C*. But Angie expresses no attitude toward that proposition, and so we are some distance from *Inconsistent Preference*. Accordingly, Gibbard doesn't analyze disagreement in attitude by appeal to logically inconsistent contents, and rather tries to motivate an intuition that some other kind of incompatibility can arise between different subjects' plans. While it is unclear whether there is any such intuitive notion, one way of trying to identify inconsistent contents is by appeal to *de se* propositions, such that every person who thinks 'If I am in *C*, I φ ' entertains the same proposition. While *de se* propositions are controversial, I'll assume them here for the sake of argument.

A different issue about inconsistency is fatal for this strategy, however. Being "in *s*'s situation" needs clarification in two respects. Does it imply *having s's preference for ends*? I'll assume it does, as necessary for accommodating Sadie's disagreement in plan with Angie.¹⁷ More importantly, does being in *s*'s situation imply *having only s's information*? To be able to account for all cases of normative disagreement as consisting in an expressed disagreement in plan, being "in *s*'s situation" must sometimes involve being in *s*'s epistemic or subjective situation, and sometimes involve being in *s*'s objective situation but with different information. Connie's postmortem statement that Angie ought to have used L and not R, for example, must express the plan to use L and not R if in *C with the same information as Angie*. But Bertie's advice that Angie ought to use R and not L must instead express the plan to use R and not L if in *C but with his own fuller information*; similarly with hindsight claims.

This scuttles any hope of explaining problematic cases of disagreement by appeal to attitudes toward inconsistent conditionals, since what I do *if I have information i_1* isn't logically related to what I do *if I have information i_2* . While this may have some welcome implications (e.g. that Connie's hindsight judgment that Angie ought to have used R doesn't disagree with her postmortem judgment that Angie ought to have used L), it also yields results incompatible with intuitions about some basic cases of normative disagreement.

In particular, it implies that better informed advisors don't genuinely disagree with the agents they advise; e.g. Bertie doesn't genuinely disagree with Angie. Bertie must express the plan to use R, *if in Angie's situation but with his own fuller information i_+* . But Angie can't have expressed the conflicting

¹⁷ This poses a problem for fundamental disagreement; see Finlay 2014, Chapter 8.

plan to use L in that very situation. Rather she must have expressed the plan to use L *in her own actual epistemic situation with information i*. The contents of these plans are logically consistent. Indeed, these plans couldn't be incompatible in any plausible sense, since a prudentially virtuous agent could and arguably should have both plans simultaneously. But Bertie's disagreement with Angie was the basic case that motivated abandoning *Inconsistent Belief* for some version of *Inconsistent Preference* in the first place. So even on its own terms the appeal to attitudes with conditional content fails to accommodate intuitions of disagreement.¹⁸

6. Conditionals with Attitudinal Content

The cases still lacking a solution involve hindsight, or better informed disagreement in a context of indirect interest, especially when the speaker doesn't prefer the end (e.g. if Sadie says 'Angie ought to have used R'). Normative claims in these contexts don't plausibly express actual preferences or similar attitudes, either toward simple or conditional contents. ERT predicts this, but it also predicts a subtly different kind of conditional attitude solution: that these normative claims pragmatically indicate *conditionals with attitudinal content*.

We've observed that the theory predicts that somebody who asserts *ought(p)* in a context of direct interest, relativized to her preferred end *e* and present information, will express the preference that *p*. So anybody who asserts *ought(p)* relativized just to her present information will pragmatically indicate the following conditional: *if her preferred end were e then she would prefer that p*. This is a counterfactual anchored in the normative belief she expressed, concerning only world-states where she still has that same belief and information.¹⁹ So Sadie's hindsight statement will pragmatically indicate that if she were to share the preference that Angie survived unharmed then she would prefer that Angie used R. The same will be true of hindsight statements where the end is preferred, like Connie's and Freddie's. I'll call these *hypothetical preferences*, in respect of their kinship to "hypothetical imperatives" and to distinguish them from the previous kind of "conditional attitudes".

Might some intuitions about normative disagreement be responsive to pragmatic expression of hypothetical preferences? A first test is to see whether directly asserting the suggested information has the same effect. Suppose the conversation had gone like this:

Angie: 'I prefer to use L.'

Sadie: 'If I were to prefer your ends, I would prefer that you use R.'

¹⁸ Gibbard indeed maintains (conversation) that there is no genuine normative disagreement from unequal information, dismissing these intuitions as mistaken. He distinguishes between the 'ought' of rationality and the 'ought' of advisability (see also Ridge 2014); however, we've also observed the need for multiple 'ought's of advisability.

¹⁹ Familiar worries about counterfactuals might be raised, e.g. that in the closest possibilities the speaker might be "irrational" so her preference for *e* isn't transmitted to *p*. The qualification 'if rational' is therefore suggested, though I believe unnecessary; see Finlay 2008, 2009, 2014.

It seems natural to say that Sadie would thereby have expressed some kind of disagreement with Angie. This is also very close to the colloquial vehicle for normative disagreement, ‘If I were you I’d φ ’. Since ‘I would φ ’ is tantamount to a description of a counterfactual preference, these utterances draw very close given that ‘If I were you...’ can be read as *if I were in your objective situation with your ends*. Whereas expressivists sometimes invoke this turn of speech to support their appeal to attitudes with conditional content, here the conditional seems to take wide scope over the attitude.

Appeal to hypothetical preferences avoids the problems observed for the expressivist’s appeal to attitudes with conditional content. First, since these conditional preferences have simple and unconditionalized contents, disagreeing claims relativized to different information (like Angie’s and Sadie’s, but also Bertie’s advice, Connie’s hindsight, etc.) thereby express hypothetical attitudes toward logically inconsistent contents (e.g. that *A uses L and not R*, and that *A uses R and not L*), so no unexplained kind of inconsistency is required. Second, the relevant attitudes can be identified simply as *preferences* about the same proposition rather than as plans or intentions, so we don’t have to countenance de se propositions or make sense of (e.g.) Sadie’s being Angie.

This also avoids a number of baroque consequences that Gibbard’s approach seems unable to evade except by positing either ambiguity or noncompositionality. We can make sense of claims that *s ought to φ* involving remote circumstances (e.g. ‘Caesar ought not to have crossed the Rubicon’) or nonagents (e.g. ‘Trees ought to have deep roots’) without saying that the speaker must have contingency plans for being in those circumstances (deliberating whether to cross the Rubicon with your legions to seize Rome, being a tree). There isn’t anything bizarre about indicating that if you were to prefer the salient end (e.g. that the Roman Republic was preserved, that trees grow strong) then you would prefer that *p* (e.g. that Caesar didn’t cross the Rubicon, that trees have deep roots).

It may seem counterintuitive that an actual disagreement could be explained by merely hypothetical attitudes. This may also seem to predict an absurd proliferation of normative disagreements. Since there are indefinitely many true counterfactuals about what somebody would prefer under some condition or other, if hypothetically preferring *p* under one set of conditions C_1 were sufficient for disagreeing with somebody who hypothetically prefers *not-p* under different conditions C_2 , wouldn’t everybody at every moment both agree and disagree with every possible normative claim?

A first response is that these hypothetical preferences and disagreements are grounded in actual expressed beliefs, although those beliefs may be logically consistent. They might rather be described as disagreements in actual dispositions, grounded in beliefs. But another necessary condition for disagreement can be adduced. Intuitively, disagreement with an attitude requires some kind of robustness, such that a disagreeing attitude *resists* the other.²⁰ These preferences differ in robustness, according to the completeness of the information on which they’re based. Angie’s

²⁰ Stevenson suggests a similar but more problematic condition: “at least one of [the disagreeing parties] has a motive for altering or calling into question the attitude of the other” (1944: 3-4). Disagreement doesn’t seem to entail any motive to change others’ attitudes.

preference to use L is provisional on no fuller information being available, for example, so it yields without resistance to Bertie’s preference based on fuller information. Bertie’s preference “trumps” hers and makes it moot.

Our pragmatic principle therefore predicts that disagreement in hypothetical preference will be asymmetrical on this basis, which I’ll call *Robust Inconsistent Hypothetical Preference*:

RIHP: B’s assertion that p normatively disagrees (agrees) with A’s assertion that q if B thereby indicates that if he were to prefer e he would on the basis of his information i_B have a preference inconsistent with (equivalent to) the preference A thereby indicates she would have on the basis of her information i_A if she were to prefer e , and $i_B \geq i_A$.

This complication is actually serendipitous. I’ve described normative disagreements asymmetrically because many are intuitively asymmetrical: Bertie intuitively disagrees with Angie, for example, but Angie doesn’t intuitively disagree with Bertie. This is observed but not explained by others,²¹ and it poses a difficulty for any theory proposing to explain normative disagreement by *Inconsistent Belief*, which implies symmetry.²² Of course, a satisfactory theory must also predict correct asymmetries, but ERT’s predictions are promising: Bertie’s advice disagrees with Angie’s deliberative claim, but not vice versa; Connie’s hindsight claim disagrees with Angie’s claim but not vice versa, and agrees with Bertie’s advice and vice versa. These results support our pragmatic account of normative disagreement, and also undermine the uncharitable hypothesis that intuitions of disagreement are simply due to blindness about the information-relativity of normative claims.

RIHP is only formulated as a sufficient condition for normative disagreement, however, and at least two obstacles prevent it from being also a necessary condition and thereby a general theory: *postmortem claims* and *negations*. The former consists in normative claims in contexts of indirect interest that are relativized to information the speaker considers inferior to her own (whether incomplete or false): for example Connie’s postmortem claim that Angie ought to have used L, and Sadie’s pseudo-advisory claim that Angie ought to use L. These won’t express hypothetical preferences as defined, since in the nearest possibilities where the speaker prefers the end they still won’t be relativized to the speaker’s full information. ERT easily accommodates the agreement expressed by these claims as agreement in belief, however, since these cases concern the same end-relational proposition Angie asserted. This suggests a disjunctive view that generates some intuitive predictions; for example, that Connie’s postmortem claim that Angie ought to have used L doesn’t disagree with Bertie’s advice or with her own hindsight claim that Angie ought to have used R, and vice versa.

Rather than simply giving a disjunctive theory, we can distinguish different kinds of normative disagreement. *Instrumental* disagreement consists in a conflict of hypothetical preferences;

²¹ Dietz 2008, Ross & Schroeder 2013.

²² An alternative hypothesis is that disagreement requires *responding* to the other claim. But cases are easily found where an earlier claim intuitively disagrees with a later one; cf. Cappelen & Hawthorne 2009, Richard 2011.

it is disagreement over what to do to achieve a particular end. *Rational* disagreement consists rather in inconsistency in end-relational belief. But this rational disagreement also implies inconsistency in a different kind of counterfactual preferences, which we can call disagreement in *doubly-hypothetical preferences*: if the speaker were to prefer that *e and only had information i*, then on the basis of her belief she would prefer that *p*. For example, Connie’s postmortem claim that Angie ought to have used L pragmatically indicates that if Connie were to prefer that Angie survived unharmed and only had Angie’s information, then she would prefer that Angie had used L. Rational disagreements can therefore also be glossed as concerning what to do to achieve a certain end with certain information.

Asserting a normative sentence can therefore express multiple attitudes. In saying ‘I ought to use L’, Angie expresses (i) an end-relational belief, (ii) an actual preference to use L, (iii) a hypothetical preference to use L on the (actual) condition that she prefers to survive unharmed, and (iv) a doubly-hypothetical preference to use L on the (actual) conditions that she prefers to survive unharmed and only has information *i*. This implies that two people may also stand in multiple relations of normative agreement or disagreement at the same time, which seems another serendipitous result: while Bertie instrumentally disagrees with Angie (in hypothetical preference), they rationally agree (in belief and doubly-hypothetical preference). He only *expresses* his instrumental disagreement, since this is the felicitous thing to do in his context of direct interest. With relevantly the same information in a context of indirect interest, Connie can rather felicitously express either her rational agreement (postmortem) or her instrumental disagreement (hindsight). Intuitions of disagreement *simpliciter* will be sensitive to whatever is the most salient kind of disagreement expressed in the context.

The second recalcitrant piece of data concerns negations, which pose a well-known problem for expressivism. Recall Freddie, whose forensic investigations revealed that the next use of R would have been fatal, and suppose now that he is informed that the round in L wasn’t a blank, causing him to believe that Angie couldn’t possibly have survived. He may then say, in hindsight:

F: ‘It was not the case that Angie ought to have used L. But not because she ought to have used R; she simply had no good options.’ (*Hindsight Negation*)

Freddie seems to disagree with Angie, but we don’t yet have an explanation for this. It can’t be a rational disagreement (in belief), because Freddie is concerned with different information. But neither does he express an inconsistent hypothetical preference, as if he said instead ‘Angie ought not to have used L,’ since *not ought(p)* doesn’t entail *ought(not-p)*. Negative claims don’t seem to express preferences at all, in which case they couldn’t express disagreement in preference.

However, our theory does predict that Freddie’s utterance indicates something about his preferences: that he *doesn’t* hypothetically prefer that Angie used L and not R.²³ He pragmatically indicates that on the basis of his belief if he were to prefer that Angie survived unharmed, he

²³ Something stronger can be said in Freddie’s case: that he is (hypothetically) *indifferent* about whether Angie uses L rather than R. This isn’t sufficiently general as an account of believing *not ought(p)*, however, since one may also do so on the basis of believing *ought(not-p)*, which we analyze as hypothetically preferring *not-p*.

wouldn't prefer that she used L.²⁴ While this doesn't identify attitudes toward inconsistent contents, it does provide a different kind of inconsistency of attitudes: preferring p and not preferring p are logically inconsistent states of mind in a broad sense, which also subsumes preferences toward inconsistent contents.

Might negations of normative claims express disagreement by virtue of this attenuated kind of "inconsistent preferences"? This runs straight into the expressivist's problem with negation. If Freddie's saying 'It wasn't the case that Angie ought to have used L' expresses merely his lack of preference for Angie's using L, how can we explain the difference with Debbie's saying 'I don't know whether you ought to use L'? Or if believing *not ought(p)* is merely lacking the preference that p , what distinguishes it from not believing *ought(p)*?²⁵ As we predict, Debbie also indicates she has no preference about Angie's using L rather than R, but unlike Freddie she doesn't intuitively disagree with Angie. Lacking an attitude doesn't in general seem to be a way of disagreeing with the attitude. But unlike expressivism, ERT identifies a significant difference between Debbie and Freddie. Whereas Debbie's lack of preference is due to lacking any relevant belief about whether Angie ought to use L, Freddie's lack of preference is grounded in his belief that it isn't the case that Angie ought to have used L.²⁶

The difference between *believing not ought(p)* and *not believing ought(p)* is therefore that whereas in one case the lack of preference results from *having* a relevant normative belief, in the other it results from *lacking* a relevant normative belief. An explanation why expressing a lack of attitude by negation can intuitively express normative disagreement then follows from our previous explanation of why inconsistency in merely hypothetical preferences can be sufficient for actual disagreement. As with those hypothetical preferences, Freddie's hypothetical lack of preference is grounded in and expressed by an actual normative belief, unlike Debbie's. The end-relational theory therefore provides a quasi-expressivist solution to the negation problem that isn't available to (pure) expressivism itself.

7. The Pragmatics of Disagreement from Different Ends

Some of our most important normative and moral disagreements are apparently fundamental, involving a conflict in basic ends.²⁷ Relational theories are accused of entailing that such disagreements are impossible, but this objection assumes some version of *Inconsistent Belief* and so no longer looks so formidable. The prospects for extending the quasi-expressivist solution to the

²⁴ What prevents Freddie from having a further desire that would cause him to prefer that A used L in the nearest possibilities? I don't have space to discuss this here, but see Finlay 2014, Chapter 8.

²⁵ E.g. Unwin 1999, 2001, Dreier 2006.

²⁶ Although Debbie believes it isn't the case that A ought *relative to D's information* to use L, in her context of lesser information this isn't relevant, and the robustness condition explains why it doesn't intuitively disagree with A's.

²⁷ Attempts have been made to accommodate the absolutist character of moral disagreement by analyzing moral claims as all relativized in the same way, and therefore as not genuinely fundamental by our definition. For discussion see Finlay 2014: Chapter 8.

disagreement problem from end-relativity look promising, since it's exactly this kind of disagreement that inspires the idea of disagreements in attitude.

Above we distinguished two different kinds of normative disagreement implied by ERT: rational disagreement in belief and doubly-hypothetical preference, and instrumental disagreement in hypothetical preference. Fundamental disagreements evidently can't be accommodated under either category, but this isn't a problem as they seem to involve a distinct kind of disagreement in any case.

Consider first contexts of direct interest, where each speaker is actively motivated toward the implicit end. (Suppose that Frank is felling a tree toward his neighbor's house, trying to get revenge for the poisoning of his dog; Dora says 'You ought to stop!', and Frank replies, 'No, I ought to do this.')

Unlike contexts of advice, here there are two different, incompatible ends. When Dora says 'You ought to stop,' (i.e. *in order not to damage your neighbor's property*, as she prefers), she thereby also expresses her preference that Frank stops. When Frank responds 'I ought to do this,' he similarly expresses his preference that he continues. Here we have attitudes toward inconsistent contents, and a disagreement in preference.

This is an expressed disagreement in *actual* rather than hypothetical preferences, which might be called *outright* normative disagreement.²⁸ Not all outright disagreements are fundamental: Angie and Bertie express preferences that are actual as well as hypothetical, for example, so Bertie disagrees with Angie not only instrumentally (about what to do to survive), but also outright (about what to do). But I maintain that all fundamental disagreements are outright. Contexts of indirect interest may seem to provide counterexamples, since here actual preferences aren't always expressed. But neither do normative claims relativized to different ends always intuitively disagree or agree with each other. Given our definition of fundamental disagreement as involving a basic conflict in preferred ends, only contexts where actual preferences are expressed (whether of direct or indirect interest) are relevant.

ERT doesn't only accommodate the existence of fundamental disagreement, but can also explain otherwise puzzling facts about its asymmetry. We've already observed one case of asymmetrical outright disagreement, between Bertie and Angie in light of the shared preference for some informational bases over others. In the case of fundamental disagreements, the analogous asymmetry is that speakers prefer some *ends* over others. Unlike the asymmetry of information, the "quality" of an end will differ between speakers, predicting less asymmetry in fundamental disagreements than in instrumental disagreements: Dora and Frank each prefers his or her own end, so their disagreement in preference is symmetrical. But other cases will be asymmetrical, wherever the second speaker makes salient an end that the first speaker is also disposed to prefer.²⁹

²⁸ To express an "actual preference" as intended here is to indicate one actually has it. One may therefore express an "actual preference" one doesn't actually have, in cases of insincerity.

²⁹ This provides a solution to an objection to relational theories made by Judy Thomson (2008) that any normative 'ought' claim—even including "hypothetical imperatives" and those concerning chess moves—can be a target of moral disagreement, and therefore must have moral content. For discussion see Finlay 2014, Chapter 8.

Whereas the problem of disagreement is widely considered fatal to relational semantics and to require either an expressivist or absolutist descriptivist account of the meaning of normative language, we've seen that the end-relational theory supports a pragmatic solution for both information- and end-relativity that accommodates ordinary intuitions of normative disagreement. Despite being semantically descriptivist, it explains these intuitions as sensitive to inconsistency in either belief or preference. While this might seem undesirably messy, we've also seen that it accommodates some data that look anomalous for nonrelational accounts despite their alleged advantages here. Normative disagreement resists analysis exclusively in terms of inconsistent beliefs and also exclusively in terms of inconsistent preferences. I conclude that the end-relational theory can only be supported, and not refuted, by our practice and sense of normative disagreement.

References

- Björnsson, Gunnar & Finlay, Stephen (2010). "Metaethical Contextualism Defended", *Ethics* 121 (1): 7-36.
- Cappelen, Herman & Hawthorne, John (2009), *Relativism and Monadic Truth*. Oxford: Oxford University Press.
- Dietz, Richard (2008). "Epistemic Modals and Correct Disagreement", in M. García-Carpintero & M. Kölbel (eds.), *Relative Truth*. Oxford: Oxford University Press.
- Dowell, Janice (2013). "Flexible Contextualism About Deontic Modals", *Inquiry* 56: 149-78.
- Dreier, Jamie (2006). "Negation for Expressivists", in R. Shafer-Landau (ed.), *Oxford Studies in Metaethics* 1. Oxford: Oxford University Press, pp. 217-33.
- Finlay, Stephen (2008). "Motivation to the Means", in D. Chan (ed.), *Moral Psychology Today: Values, Rational Choice, and the Will*. Springer, pp. 173-91.
- Finlay, Stephen (2009). "Against All Reason? Skepticism About the Instrumental Norm", in C. Pigden (ed.), *Hume on Motivation and Virtue*. Palgrave MacMillan.
- Finlay, Stephen (2014). *Confusion of Tongues: A Theory of Normative Language*. New York: Oxford University Press.
- Gibbard, Allan (2003). *Thinking How to Live*. Cambridge, MA: Harvard University Press.
- Horgan, Terence & Timmons, Mark (1992). "Troubles on Moral Twin Earth: Moral Queerness Revived", *Synthese* 92 (2): 221-260.
- Kölbel, Max (2002). *Truth Without Objectivity*. London: Routledge.
- Kolodny, Niko & MacFarlane, John (ms). "Ought: Between Subjective and Objective". Unpublished manuscript.
- MacFarlane, John (2007). "Relativism and Disagreement", *Philosophical Studies* 132: 17-31.
- Regan, Donald (1980). *Utilitarianism and Cooperation*. Oxford: Oxford University Press.
- Richard, Mark (2011). "Relativistic Content and Disagreement", *Philosophical Studies* 156 (3): 421-31.
- Ridge, Michael (2014). *Impassioned Belief*. Oxford: Oxford University Press.

- Robinson, Denis (2009). "Moral Functionalism, Ethical Quasi-Relativism, and the Canberra Plan", in D. Braddon-Mitchell & R. Nola (eds.), *Conceptual Analysis and Philosophical Naturalism*. Cambridge, MA: MIT Press.
- Ross, Jacob & Schroeder, Mark (2013), "Reversibility and Disagreement", *Mind* 122: 43-84.
- Stevenson, Charles L. (1944). *Ethics and Language*. New Haven, CT: Yale University Press.
- Thomson, Judith Jarvis (2008). *Normativity*. Chicago, IL: Open Court.
- Unwin, Nicholas (1999). "Quasi-Realism, Negation and the Frege-Geach Problem", *Philosophical Quarterly* 50 (196): 337-52.
- Unwin, Nicholas (2001). "Norms and Negation: A Problem for Gibbard's Logic", *Philosophical Quarterly* 51: 60-75.