

Why Plan-Expressivists Can't Pick up the Moral Slack

Margaret Shea

1 Introduction

Good moral deliberation often involves taking into account what ought to be done, but likely will not be done. Indeed, a certain form of reasoning is common and cogent, both in the interpersonal and intrapersonal case:

- (i) B should φ , but he won't – so A should φ .
- (ii) C should φ , but C won't – so C should ψ .

It is a familiar fact that our normative judgments about what an agent should do not disappear once it is stipulated that the agent will (or will likely) not do that thing. Consider the following cases:

LONELY AUNT: Ann's brother Bob has promised to visit their lonely great-aunt at the nursing home this weekend. But he is selfish and unreliable, and Ann knows he's booked a luxury getaway. Due to COVID restrictions, the great-aunt can only have one visitor per weekend. Bob should visit, but he won't – so Ann should.

SICK UNCLE: Claire's uncle is in the hospital and needs company. Claire hates the smell and ambience of hospitals, and always fails to visit when she says she will. Whenever she drives to the hospital, she gets drawn in by the spa next door, and goes for a massage instead of visiting. Claire can either drive to the hospital and visit her uncle in person, or stay home and call him, but cannot do both. Visiting in person is better than talking on the phone. Claire should get in the car and visit, but she won't – so she should stay home and call.¹

We judge that Bob should visit his great-aunt and that Claire should visit her uncle. In each case, it is foreseeable that an agent will fail to do something. This is perfectly consistent with our judging that she *should* do it.²

¹ Worsnip (2021: 214) discusses a similar case.

² Judgments about what an agent "should" or "ought to" do refer to normative judgments about what an agent has all-things-considered reason to do. I use "should" and "ought" interchangeably.

These cases involve what I will call “non-corealizable” actions.³ A pair of actions (by a single individual, or by different individuals) is non-corealizable if and only if at most one of them can be done. You could win the race, and so could I, but we cannot both win. So your winning and my winning are non-corealizable. Let us call a pair of actions with this feature “interpersonally non-corealizable,”⁴ and a pair of actions “intrapersonally non-corealizable” when a single agent can φ , and can ψ , but cannot do both. I can go to Cancun or Paris for spring break, but I cannot go to Cancun and Paris. These actions are intrapersonally non-corealizable.

What these two cases – LONELY AUNT and SICK UNCLE – seem to suggest is that “should”-judgments can be jointly coherent even when they jointly recommend non-corealizable actions. The judgments that A should φ and B should φ can be jointly coherent, even though it is not possible for both A and B to φ . Similarly, the judgments that C should φ and should ψ can be jointly coherent, even though it is not possible for C both to φ and to ψ . But if this is correct, then many expressivist theories, including the *locus classicus* of plan expressivism (Gibbard 2003), are in trouble. Plan-expressivists associate normative judgment with an attitude that satisfies a corealizability constraint. The result is the introduction of a consistency requirement on normative judgments where none is warranted – or, more cautiously vis-à-vis Gibbard, where one is warranted only by the introduction of a contentious first-order premise, possibilism.⁵

Non-corealizability raises a second problem for plan-expressivism, elaborated in §4. To anticipate: while standard plan-expressivism (Gibbard 2003) associates every normative judgment with

³ I adapt the term from Ayars (2022).

⁴ Nothing hinges on the fact that, in my examples, it is one and the same act-type, φ , which each of two agents should perform, but which they cannot both perform. The problem I raise clearly generalizes to cases involving two act-types, A’s φ -ing and B’s ψ -ing, which each should be performed, but cannot simultaneously be performed.

⁵ Compare Van Roojen (1996: 322): “We have noted a recurring problem [for expressivists], that the resources used to explain inconsistency where all agree there is inconsistency generate inconsistency where there seems to be none.”

a *planning subject* (defined below), cases involving interpersonally non-corealizable acts prompt normative judgments which are not plausibly associated with any such subject.

Allan Gibbard's 2003⁶ and Alisabeth Ayars' 2022⁷ views serve as exemplars of the plan-expressivist program, and I focus on them in what follows. But the objections I raise apply to any expressivist view which identifies normative judgment with an attitude that is subject to a corealizability constraint. On Gibbard's (2003) account, normative judgments are planning states. A "plan" is a desire-like attitude which "[determines] what to do in various contingencies, expected or hypothetical";⁸ it is that state which settles what to do, concluding deliberation and characteristically issuing in intentions.⁹ Gibbard holds that first-personal normative judgments of the form "I should φ " amount to a plan, on my part, to φ . Third-personal normative judgments, by contrast, are contingency plans for the counterfactual scenario in which I, the judge, am also the subject of the judgment (hereafter: the "planning subject"). That is: the judgment I express with "Jane should φ " amounts to a plan, on my part, to φ if I am in Jane's exact circumstances.¹⁰

Ayars' (2022) expressivist account of normative judgment follows Gibbard's in most respects. Normative judgments are associated with an intention-like, non-cognitive attitude called "decision." For the purposes of this paper, an Ayars-style "decision" can be understood as a Gibbard-style "plan."¹¹ The novelty of Ayars' view lies in the fact that it gives a uniform analysis of first and third-

⁶ Gibbard (2003).

⁷ Ayars (2022). Ayars (this volume) briefly summarizes Ayars (2022).

⁸ Gibbard (2003: 53).

⁹ I borrow this helpful characterization from Ayars (2022: 10-11).

¹⁰ Gibbard defends an indexical theory of belief, desire, and intention-states. So on his account, all plans are *de se*. When I make a plan for what to do if I am in Jane's circumstances, I am making a hypothetical plan for what to do in that set of worlds whose center is Jane, where Jane's descriptive circumstances are those she is actually in.

¹¹ In fact, Ayars (2022) more than Gibbard insists on a distinction between plans and future intentions. One might form a future intention always to go left in a Buridan's Ass case, but a "plan" of this sort is not a "decision" in Ayars' sense. This seems the right result: someone who makes such a plan, does not thereby judge that she *should* go left.

personal normative judgments. All normative judgments of the form “ X should φ ” amount to *decisions that* $X \varphi$. A first-personal normative judgment is a decision I make for myself. A third-personal normative judgment is a decision I make for another. My judgment that I should φ is a decision to φ . My judgment that Jane should φ is a decision, on my part, that Jane φ .¹²

The variations in Ayars and Gibbard’s views affect how the problems raised by non-corealizability manifest themselves. But non-corealizability presents a stumbling block for both. Let us first examine a problem which interpersonal non-corealizability raises for Ayars’ expressivism.

2 The Challenge from Interpersonal Non-corealizability

2.1 The Challenge

On Ayars’ view, intrapersonal and interpersonal normative judgments are associated with a single attitude that is subject to a consistency constraint. Again, all normative judgments of the form ‘ X should φ ’ amount to decisions that $X \varphi$. Decisions are subject to a corealizability constraint: an agent cannot coherently decide both to shave her head and to put rollers in her hair. It is therefore incoherent to judge that an agent should (now) shave her head and should (now) put rollers in her hair. More generally, it is incoherent to decide that $A \varphi$ and that $B \psi$ if one believes that’s A ’s φ -ing and B ’s ψ -ing preclude one another. Since normative judgments are decisions, one can coherently judge that A should φ and that B should ψ only if one believes that it is possible for A to φ when $B \psi$ ’s.¹³

Now recall the sad case of LONELY AUNT. We judge that Bob should visit, but we foresee that he won’t. So, we judge, Ann should visit. According to Ayars, these judgments are decisions. The

¹² See Ayars (2022) for exposition and defense of her idea that one agent can “decide that” another do something.

¹³ As Ayars (2022, this volume) emphasizes. The interpersonal co-realizability constraint on normative judgment plays a crucial role in her arguments that rational egoism is incoherent (2022: 16-19), and that practical reasons are essentially public (this volume). My objection thus targets the very feature of Ayars’ view which she relies upon in defending various normative claims.

decision that Bob visit, and the decision that Ann visit, cannot both be realized. Our pair of judgments is made in full awareness of this fact. So Ayars' view must say that the pair of judgments is incoherent. But that is false. The judgment that Ann should visit her aunt puts no rational pressure on the judgment that Bob should visit her – as is obvious if we consider the situation from Ann's perspective. Upon resolving to pick up her brother's slack, Ann's judgment that he ought to be doing the job remains intact. The judgment she expresses by means of 'I should visit' does not put rational pressure on her to jettison her judgment, 'Bob should visit.'¹⁴ He does not get off the hook so easily.

2.2 *Objection to the Interpersonal Challenge*

Ayars' view fails only if there are genuine cases about which it is coherent to judge that each of two agents should φ , although (as one believes) they cannot both φ . Using LONELY AUNT as an example, let me address two objections to the claim that this can in fact be coherent. First, one might object that it is not right to say without qualification that Ann should visit. That judgment is underdescribed, and implicitly includes a proviso: 'Ann should visit *if* Bob does not visit.' If *that* conditional thing is what we intend when we say that Ann should visit, the objection goes, then there is no problem for Ayars' view. For, once this imprecision is eliminated, our pair of judgments ('Bob should visit,' 'Ann should visit if Bob does not') expresses a pair of decisions that are corealisable. It is possible both for Bob to visit, and for Ann to visit if Bob does not. So Ayars's view does not deem *this* pair of judgments incoherent.

The objection fails, however. Suppose that – as the objection proposes, and as is apt – we judge 'Bob should visit' and 'Ann should visit if Bob does not.' Then we learn that Bob will not visit.

¹⁴ Since Ayars' view associates intrapersonal and interpersonal judgments with a single attitude, we are free to redescribe the case, and the relevant judgments, from Ann's perspective. The distinction between (i) the pair made up of the third-personal judgments "A should φ " and "B should φ " and (ii) the pair made up of the first-personal judgment "I should φ " and the third-personal judgment "B should φ " comes to nothing, because sets of "should" judgments get the same analysis, and are subject to the same corealizability constraint, regardless of who is making them.

This should lead us to form the unconditional judgment that Ann should visit, while *retaining* the judgment that Bob should visit. But these judgments are incoherent on Ayars' view, for they express non-corealizable decisions. Appealing to possible worlds sharpens the point. The possible world which concerns us, and presents a problem for Ayars, is the one in which Bob fails to visit. *Contra* Ayars, it is coherent to judge that, in this world, Ann should visit, and that, in this same world, Bob should visit.¹⁵ Bob's foreseeable failure to do what he should does not make it false that he should do it. That is why Ann can cogently construe her fulfillment of her obligation as *picking up* Bob's slack.

The same consideration undermines a broader version of this objection, namely that 'Ann should visit' and 'Bob should visit' express decisions for mutually exclusive sets of descriptive circumstances. On the plan-expressivist story, normative judgments are always plans (or, in Ayars' lingo, decisions) *for* some descriptive circumstances c (the 'fact' half of a 'fact-prac world' (Gibbard (2003: 47))). ' X should φ in c ' expresses the decision that X φ if the circumstances (set of descriptive facts) c obtains. The proposal is that 'Bob should visit' expresses a plan for circumstances c_1 , while 'Ann should visit' expresses a plan for some distinct circumstances c_2 . So long as c_1 and c_2 preclude one another (they cannot both obtain), these judgments express corealizable decisions, and are coherent.¹⁶

This objection, too, fails because Bob *should* visit even in circumstances where he does not. It is coherent to judge that Bob should visit in some circumstances c and that, *in these same circumstances* c , Ann should visit.¹⁷ Define c as the exhaustive set of descriptive facts about the world up to and

¹⁵ Consider that if neither visits, it is appropriate for *both* to feel guilty.

¹⁶ A 'decision' in Ayars' sense counts as realized when its condition goes unsatisfied. Therefore, if c_1 and c_2 preclude one another, the decisions 'Bob visit if c_1 obtains' and 'Ann visit if c_2 obtains' are corealizable.

¹⁷ A subtler proposal is that there is a context-shift across our judgments about Ann and Bob. (Mellor & Shea (ms.) supports this suggestion.) According to Kratzerian contextualist semantics for 'ought,' 'A ought to φ ' is true relative to some modal base (set of propositions held fixed by a speaker) m , if and only if, in the best worlds consistent with m , A φ 's (Kratzer (2012), Worsnip (2019)). Contextualists will point out that there is no single modal base relative to which 'Ann should visit' and 'Bob should visit' are true: 'Ann should visit' is true only relative to a modal base which includes the proposition that Bob will fail to visit, while 'Bob should visit' is false relative to any modal base which includes this proposition (since there are no worlds compatible with this proposition in which Bob visits).

including the time at which Bob boards the flight to his getaway. Ann should visit in c because Bob is going to fail to do so. And Bob should visit in c because he promised that he would. Again, he does not escape his obligation by foreseeably failing to fulfill it.

3 The Challenge from Intrapersonal Non-corealizability

3.1 The Challenge

Gibbard avoids Ayars' unfortunate result in cases involving interpersonally non-corealizable actions – *viz.* the result that it is incoherent to judge that an agent A should φ and that another agent B should φ , whenever one believes A's and B's φ -ing to be incompatible.¹⁸ On Gibbard's account, a third-personal judgment is a contingency plan for what to do in another agent's shoes.¹⁹ One's judgments that A should φ and that B should φ express two plans for two different scenarios: one, a plan to φ if one finds oneself in A's descriptive circumstances; and another, to φ if one finds oneself in B's. These two plans are corealizable, even if A's φ -ing and B's φ -ing preclude one another. For one will not find oneself in both of A and B's descriptive circumstances at the same time, and conditional plans (*e.g.* 'if in A's circumstances, φ ') count as realized when the relevant condition goes unsatisfied. Thus, on Gibbard's account, the fact that it is impossible for A to φ whenever B φ 's has no bearing on the coherence of judging that A should φ and B should φ .²⁰ Take LONELY AUNT. On Gibbard's

Could Ayars' expressivism incorporate Kratzerian contextualism, and avoid the challenge? It is not clear that it could, since Kratzerian contextualism gives a descriptivist analysis of 'ought.' Supposing that it could, it is not clear how doing so would address the challenge I have raised. The imagined contextualist-expressivist theory would have to introduce some subtler analysis of 'ought' and some subtler coherence constraint on normative judgment than Ayars provides. For *ex hypothesi*, the semantic values of 'ought' judgments vary as the modal base varies, while Ayars (2022, this volume) is silent on how this multiplicity of semantic values could be understood in terms of decisions. In the absence of a developed hybrid view, which answers this question and identifies a subtler coherence constraint on normative judgment, it is difficult to say whether accepting 'ought' contextualism would help Ayars' view address my objection. Even if it would, contextualism is a costly toll to pay for a plan-expressivist view.

¹⁸ See note 4.

¹⁹ See note 10.

²⁰ *Mutatis mutandis* for the pair of judgments 'I should φ ' and 'B should φ ,' where my φ -ing and B's φ -ing are interpersonally non-corealizable. The former judgment amounts to a plan, on my part, to φ in my actual circumstances;

view, the judgment that Bob should visit is a plan to visit if in Bob's exact descriptive circumstances. The judgment that Ann should visit is a plan to visit if in Ann's exact descriptive circumstances. One cannot be in both of Ann and Bob's exact descriptive circumstances at the same time. So the fact that it is impossible for Ann to visit whenever Bob visits has no bearing on the coherence of judging that Ann should visit and Bob should visit.²¹

However, the problem I now raise concerning intrapersonal non-corealizability affects Ayars and Gibbard equally. Since they face the exact same quandary, for the exact same reason, I will lump them together in the remainder of this section under the name "plan-expressivists," and talk of "plans" rather of "decisions." It is incoherent to plan to perform each of two acts which one believes to be mutually exclusive in a given circumstance. So much is uncontroversial. According to plan-expressivists, intrapersonal normative judgments – judgments about what a single agent should do – just are plans for that agent. So intrapersonal normative judgments are subject to this same constraint: it is incoherent to judge that an agent should perform each of two acts which one believes to be mutually exclusive in a given circumstance.

Yet, just as one's judgments that A and B should φ can be jointly coherent even when one does not believe that A and B both can φ , so one's judgments that a single agent A should φ and should ψ can be jointly coherent, even when one does not believe that A can both φ and ψ . This much plan-expressivists must allow, unless they assume a contentious normative premise, possibilism. Recall SICK UNCLE: Claire can either get in the car and visit her uncle in the hospital (φ) – which would be best, but which foreseeably she will fail to do – or she can stay at home and telephone him (ψ). She cannot do both. (Suppose that patients have only one hour for socializing a week, and that there is no

the latter amounts to a plan to φ in Bob's circumstances. My actual circumstances and Bob's circumstances are distinct, so my plans are corealizible, and my judgments coherent.

²¹ Thanks to Jamie Dreier for discussion of this point.

cellphone signal within 100 miles of the hospital, such that driving to the hospital precludes calling.) Note that whether Claire successfully gets in the car and visits her uncle is under her agential control. Though it is foreseeable that she will fail, it is not *impossible* for her to succeed, and if she fails, this will be a failure for which she is responsible. Addiction and compulsion undermine genuine agency, but it is not as though Claire is addicted to the seductive eucalyptus scent wafting from the spa's entrance, or compulsively averse to entering the hospital.

It is uncontroversial that, all things considered, Claire should drive to the hospital and visit her uncle in person.²² But is it also true that, all things considered, Claire should stay home and call her uncle? The answer to this question depends on whether “actualism” or “possibilism” is correct. These are rival accounts of the appropriate frame of comparison to adopt in determining whether an agent should perform act or act-sequence χ . Possibilism is the view that an agent should χ if and only if, of all the act-sequences available to her, χ -ing is, or is part of, the best act-sequence (that is: if the best world that she can bring about is a world in which she χ 's).²³ Actualism is the view that an agent should χ if and only if her χ -ing is better than her not χ -ing (that is: if the nearest world in which she χ 's is better than the nearest world in which she does not χ).²⁴ Both views have considerable intuitive appeal and counterintuitive consequences.²⁵ Which to accept is a substantive normative question, the resolution of which is, for the purposes of this argument, beside the point. What matters is that possibilism and actualism entail different judgments about Claire's situation and, more precisely, that

²² Actualists and possibilists ground this judgment in different considerations, described below.

²³ This formulation is adapted from Timmerman & Cohen (2016).

²⁴ One may prefer to formulate the distinction in terms of expected value. Possibilism is the view that an agent should χ if and only if, of all the act-sequences available to her, χ -ing is, or is part of, the act-sequence with the highest expected value. Actualism is the view that an agent should χ if and only if the expected value of χ -ing is higher than the expected value of not- χ -ing.

²⁵ Timmerman & Cohen (2016).

only possibilism supports the introduction of an intrapersonal non-corealizability constraint on normative judgment.

Possibilists hold that Claire should drive to the hospital and visit in person, on the grounds that this is the best act-sequence available to her. They deny that Claire should stay home and call, on the grounds that this is not the best act-sequence available to her. More generally, possibilism guarantees that whenever φ and ψ are intrapersonally non-corealizable, one of the two judgments ‘C should φ ’ and ‘C should ψ ’ is mistaken. Possibilism thus warrants an intrapersonal corealizability constraint on normative judgment. Actualists agree with possibilists that Claire should get in the car and visit, but they explain this fact differently. Claire should get in the car and visit because it is better that she do so than that she not do so. But for the same reason, Claire should stay home and call: it is better that she stay home and call than that she not do so.²⁶ So, actualists say, Claire should leave the house and visit her uncle; and she should stay at home and telephone him.

Of course, there is no world in which Claire can live out both of the plans which these judgments express. So, the plan-expressivist must say, the actualist’s judgments are incoherent. But the imposition of this coherence constraint is appropriate only on the assumption that possibilism is true. This is a doubly bad result for the plan-expressivist. First, many people find possibilism

²⁶ According to the actualist, it is better that Claire stay home and call her uncle than that she not do so because, in the nearest possible world in which she does not do so, she does not drive to the hospital and visit; her uncle is altogether ignored.

implausible.²⁷ Second, most expressivists do not want the viability of their meta-ethical theory to hinge on a contentious answer to a controversial normative question.²⁸

3.2 *Objections to the Intrapersonal Challenge*

It may be objected that, in cases like SICK UNCLE, actualists are not committed to two normative judgments which express non-corealizable plans. Actualists must hold that Claire ought to get in the car and visit, and that Claire ought to stay home and call. But it may be proposed that only one of these judgments is *action-guiding*, and so analyzable as a plan.²⁹ *Contra* this proposal, *prima facie* both actualist ‘ought’ judgments are of the kind standard plan-expressivism sets out to analyze: *viz.*, action-guiding, all-things-considered ‘ought’ judgments, which predicate “to-be-doneness” of actions (Gibbard 2003: 13).³⁰ Standard actualism holds not that one ‘ought’ judgment about Claire is action-guiding and the other not, but that both are action-guiding judgments, which hold relative to distinct option-sets. After all, actualists adduce pairs of judgments like ‘Claire ought to get in the car and visit’

²⁷ See Timmerman & Cohen (2016). Possibilism struggles to do justice to the idea that agents ought to plan around their anticipated moral failings. Imagine a person who always blows up at colleagues in faculty meetings. Although it would be best for him to attend faculty meetings and manage his temper, and although this is possible (losing his temper is under his control), it is plausible that, after repeated outbursts, he is under an obligation to stop placing himself in the temper-triggering situation. Possibilism struggles to explain this because it enjoins agents always to act as the best act-sequence requires, even when there is overwhelming evidence that they will fail to execute it. See Kiesewetter (2015, 2018) for a reply on the possibilist’s behalf, and White (2017) for criticism of Kiesewetter (2015).

²⁸ In a primer on moral cognitivism versus non-cognitivism, Mark Van Roojen (2018) writes: “Many think it a desideratum in metaethical theorizing that a candidate theory be consistent with all or most normative theories actually defended by serious normative ethical proponents.” Compare remarks in Gibbard (1990: 88) and Van Roojen (1996: 324) regarding the relation between the possibility of moral dilemmas and proposed solutions to the embedding problem (also cited in Van Roojen (2018)).

Whether to accept possibilism or actualism is a normative ethical question. The choice is between two accounts of what our normative reasons or ‘ought’s consist in. Theories about what our normative reasons or ‘ought’s consist in are normative ethical theories, rather than meta-ethical theories, on any conventional understanding of that distinction. (Cf. Estlund 2019: 166-167)).

²⁹ A related proposal is that the actualist’s two judgments about Claire feature semantically distinct ‘ought’s. If the proposal is that both ‘ought’s are action-guiding, then it is a version of the proposal addressed below (p. 14). If the proposal is that only one is action-guiding, then the proposed view is not conventionally actualist (n. 31). Thanks to Alex Worsnip for discussion.

³⁰ “*Ought* questions...are by their very nature questions of what to do. Understanding this is the way to understand what *ought* assertions mean...[*ought* claims] are claims about what to do:” Gibbard 2003: 9-10.

and ‘Claire ought to stay home and call’ as *evidence that* the action-guiding ‘ought’ is always relativized to an option-set.³¹ A version of plan-expressivism which denies that both judgments feature the action-guiding ‘ought’ is to that extent less actualist.

Moreover, the objection is undermotivated. The main reason to think that only one of the actualist’s judgments about Claire is action-guiding is a principle I will call ENKRATIC HARMONY:³² an agent cannot be rationally committed to an incoherent intention-set. Benjamin Kiesewetter (2015: 931), *e.g.*, appeals to this principle in arguing that that agents cannot have conclusive deliberative reason to perform two non-corealizable acts. Stipulate that Claire is a perfectly informed actualist. She judges that she ought to get in the car and visit, and that she ought to stay home and call. Then – granting the widespread assumption that agents are rationally committed to intend in accord with their action-guiding ‘ought’s – ENKRATIC HARMONY is violated. So it must be false that Claire deliberatively ought to get in the car and visit, and ought to stay home and call.

³¹ To elaborate: when answering open-ended ‘ought’ questions of the form ‘What ought Claire now to do?’, actualism directs us to consider “all the maximally relevantly specific actions possible at or during that time.” When answering action-specific ‘ought’ questions of the form ‘Ought Claire to φ ?’, actualism directs us to consider “the set consisting of the action and what the agent would do instead” (Jackson & Pargetter 1986: 255). So, granting actualism, the answer to the question ‘What ought Claire to do?’ is ‘Get in the car and visit;’ the answer to the question ‘Ought Claire to stay home and call?’ is ‘Yes;’ and the answer to the question ‘Ought Claire to get in the car and visit?’ is also ‘Yes.’ The relevant point for our purposes is that actualists do not distinguish between the *kind* of ‘ought’ which figures in ‘Claire ought to stay home and call’ and ‘Claire ought to get in the car and visit.’ Both are taken to be instances of the conclusive, deliberative ‘ought.’ Notably, this is precisely the feature of actualism which Kiesewetter (2015) attacks. My interest is in observing that it is a part of the seminal statement of actualism (Jackson & Pargetter 1986), not in defending it.

³² Kiesewetter (2015: 930-934) offers a second reason to think that only one of the actualist’s judgments about Claire is deliberative. He holds that an ‘ought’ judgment is deliberative only if it provides a ‘rational basis’ for making a decision about what to do, and that it is not rational to make a decision “on the basis of a judgment that an action is recommended out of a subset of options.” Since Claire’s judgment ‘I ought to stay home and call’ recommends this act-sequence out of a subset of her options, it is not deliberative. I cannot address Kiesewetter’s argument in detail, but offer two brief replies. First, actualists should reject Kiesewetter’s claim that the *only* rational basis for making a decision is a judgment that an action is recommended out of all options. Actualists face the burden of identifying a procedure for rational decision-making in situations like Claire’s, but in doing so they need not give up the claim that Claire’s judgments ‘I ought to get in the car and visit’ and ‘I ought to stay home and call’ are both deliberative. For, second, although there is indisputably a tight connection between deliberative judgment and practical deliberation, this does not entail Kiesewetter’s definitional claim, that a judgment is deliberative only if it “settles the conclusion of practical deliberation,” where this means that it recommends an action not out of some subset of options, but out of all options. Plausibly, deliberative ‘ought’ judgments *characteristically and generically* play this role, but are obstructed from doing so in cases like Claire’s, where an agent is irrational. Kiesewetter’s definition of deliberative judgment strikes me as too strong, and difficult to motivate without assuming possibilism.

This is too quick. ENKRATIC HARMONY looks most plausible *on the assumption* that agents are never in situations in which they deliberatively ought to perform each of a set of non-corealizable actions. But that is just what actualism denies. Moreover, as Stephen White (2017: 728) observes, to avoid the implausible result that “*rationality itself* [could] require an incoherent set of intentions” (my emphasis), it is not necessary to claim that agents cannot be rationally committed to an incoherent intention-set. It is enough to claim that, whenever an agent is thus committed, this fact traces back to her irrationality. Cases like Claire’s are ones in which an agent is, by stipulation, irrational (*e.g.* foreseeably akratic). *Provided that she is irrational*, White argues, an agent’s ‘ought’s can rationally commit her to an incoherent intention-set. I am persuaded by White’s argument.³³ But even those who are not persuaded should see that plan-expressivism is in trouble. As White and Kieseewetter’s exchange attests, ENKRATIC HARMONY occasions live normative-ethical disagreement. So whether it is possibilism, or the principle that an agent cannot be rationally committed to an incoherent intention-set, which warrants the introduction of a corealizability constraint on normative judgment, the problem facing plan-expressivism is the same. That constraint is warranted only by the introduction of a contentious normative premise.

A distinct objection is that ‘Claire ought to get in the car and visit’ and ‘Claire ought to stay home and call’ express plans for two mutually exclusive sets of descriptive circumstances. This is the analogue of an objection raised, above, to my discussion of LONELY AUNT (§2.2). Recall that normative judgments are plans *for* some descriptive circumstances, and that conditional plans count as realized when the relevant condition goes unsatisfied. If the actualist’s judgments about Claire are that she should get in the car and visit in circumstances c_1 and that she should stay home and call in

³³ But *see* Kieseewetter’s (2018) response to White’s (2017) response to Kieseewetter (2015).

some mutually exclusive circumstances c_2 , then these judgments express corealizable plans, and are coherent.

The objection fails because, granting actualism, ‘Claire ought to get in the car and visit’ and ‘Claire ought to stay home and call’ must express plans for a common set of descriptive circumstances. Define c as the fully specific set of propositions about the way the world is at the time at which Claire has to decide whether or not to drive to the hospital. Actualism entails that *in* c Claire ought to get in the car and visit her uncle, *and* ought to stay home and call him. Thus, actualism entails a pair of judgments which express non-corealizable plans.³⁴

Of course, plan-expressivists drawn to actualism could modify their theory in light of the challenge I have raised. Standard plan-expressivism holds that normative judgments are plans for some particular agent X in some particular descriptive circumstances c . An actualist-friendly plan-expressivism would have to add that these plans are relative to the options X has in c , and that there are two admissible construals of X ’s options in c , such that ‘ought’ judgments are ambiguous between plans for two different option-sets. On the imagined view, the corealizability constraint on normative judgment is internal to each disambiguated class of plans: when choosing from *one and the same set of options*, it is incoherent to judge that an agent should perform each of two acts which one believes to be mutually exclusive in a given circumstance. The proposal will then be that ‘Claire ought to get in the car and visit’ expresses a plan relativized to one option-set while ‘Claire ought to stay home and

³⁴ Just as they discern a context-shift across our judgments about Ann and Bob (note 17), Kratzerian contextualists will discern a context-shift across the actualist’s two judgments about Claire. There is no modal base relative to which ‘Claire ought to get in the car and visit’ and ‘Claire ought to stay home and call’ are both true. For the latter is true only given the proposition that Claire will fail to get in the car and visit, while the former is false given this proposition (there is *no* world compatible with it in which Claire gets in the car and visits, so it cannot be that in the *best* world compatible with it, Claire gets in the car and visits). As aforementioned, is not clear that plan-expressivists can adopt Kratzerian ‘ought’ contextualism, given its descriptivist analysis of ‘ought.’ Assuming they can, it is not obvious what follows from the observation that the actualist’s judgments about Claire are made relative to different contexts. I suspect it supports a version of plan-expressivism on which plans are relativized to option-sets (p. 14). Thanks to Alex Worsnip for discussion of this point. Worsnip (2021: 213-15) offers a contextualist take on cases like SICK UNCLE.

call' expresses a plan relativized to the other (which excludes the option of getting in the car and visiting), with the result that the actualist's pair of judgments are coherent.

I take no position on whether plan-expressivists should modify their view in this way. Whether they do or not, my main point stands. To introduce a distinction between (familiar) plan-expressivism, which precludes actualism, and (revised) plan expressivism, which permits it, is to concede that in developing their account of what planning is, plan-expressivists will have to choose between actualism and possibilism. For this is what will determine, first, whether there is an additional (ambiguous) parameter – options – relative to which all plans are made, and second, what scope the corealizability constraint on normative judgment has.

4 Conjunctive Subjects: Unfinished Business for the Plan-Expressivist

4.1 Normative “Clumping”

The cases which looked to make trouble only for Ayars – cases in which it is coherent to judge that each of two agents should φ , although it is impossible for both to φ – draw our attention to a broader phenomenon, which (among numerous intriguing ramifications) threatens the standard plan-expressivist analysis of third-personal normative judgments, associated with Gibbard. Cases like LONELY AUNT show there to be a distinction between (i) the conjunction of the judgments that A should φ and that B should φ and (ii) the judgment that φ -ing should be done *jointly* – i.e. *by the pair* of A and B. When pointed out, this distinction may seem obvious, but it is easily overlooked: ordinary language obscures it,³⁵ standard deontic logic ignores it, and in many cases, it is normatively insignificant.

³⁵ Some may deny that ‘A and B should φ ’ is ambiguous, in ordinary language, between two interpretations. Some English speakers report that the *only* interpretation of this sentence they hear is one on which it expresses the judgment that A and B jointly should φ . Thus they deny that, in LONELY AUNT, there is any interpretation of ‘Ann and Bob should visit’ which is true. *Contra* these reports, it seems appropriate to say ‘Ann and Bob should feel guilty that their aunt feels lonely; after all, they should have visited.’ But nothing much hinges on this. For it is *Gibbard’s view* which assumes that ‘A and B should φ ’ is semantically equivalent to the conjunction of ‘A should φ ’ and ‘B should φ .’ This is the feature of the

‘A and B should φ ’ has two semantically distinct interpretations. So I will introduce some disambiguating terminology.

- (i) The judgment ‘each of A and B should φ ’ is written ‘A and B_e should φ ,’ and is simply the conjunction of ‘A should φ ’ and ‘B should φ .’³⁶
- (ii) The judgment ‘A and B jointly should φ ’ is written ‘A and B_j should φ ,’ and holds if and only if φ -ing is a thing which should be done *by the pair* of A and B.³⁷
- (iii) *Clumping* refers to the phenomenon wherein the conjunction of ‘A should φ ’ and ‘B should φ ’ (*i.e.*, ‘A and B_e should φ ’), together with relevant features of the context, grounds the further judgment ‘A and B_j should φ .’³⁸

Cases like LONELY AUNT show that sets of “should” judgments do not always clump. The judgments that A should do some thing and B should do this same thing, *i.e.* the judgment that A and B_e should do this thing, does not entail the judgment that they jointly should do that thing. For is not always the case that the normative considerations which support the performance of an act by each of a set of individuals, support the performance of this act by all those individuals. So, not all pairs of atomic judgments (judgments about individual agents) compose into a judgment about what this pair of agents should do. ‘A should φ and B should φ ’ and ‘A and B jointly should φ ’ have distinct truth-conditions.³⁹

view I attack. If one can felicitously assert ‘A and B should φ ’ only when one judges that A and B jointly should φ , so much the worse for Gibbard. He does not provide any bespoke analysis for judgments of this form, which – I argue – are not semantically equivalent to the conjunction of ‘A should φ ’ and ‘B should φ .’

³⁶ *E.g.* in LONELY AUNT, ‘Ann and Bob_e should visit’ is true.

³⁷ *E.g.* in LONELY AUNT, ‘Ann and Bob_j should visit’ is false. The word ‘joint’ is not meant to evoke teamwork.

³⁸ *E.g.* in LONELY AUNT, ‘Ann should visit’ and ‘Bob should visit’ do not clump: the judgment ‘Ann and Bob_e should visit’ does not entail the judgment ‘Ann and Bob_j should visit.’ Clumping is similar to agglomeration. I refrain from using that term because the cases under consideration involve multiple necessity operators (“Ann-ought,” “Bob-ought,” *etc.*), while agglomeration involves only one.

³⁹ More generally, there is a semantic difference between the conjunction of atomic judgments ‘A ought to φ and B ought to ψ ’ and the judgment ‘A and B [jointly] ought to φ and ψ [respectively]’ (Mellor & Shea (ms)).

This fact is easily overlooked, perhaps because clumping failures are relatively rare. With the right background conditions, unclumped judgments (of the form ‘A and B_e should φ ’) ground their clumped counterparts (of the form ‘A and B_j should φ ’). For example, the judgment ‘Trump and Barr should resign’ is partly grounded in the pair of judgments that Trump should resign and that Barr should resign (*i.e.* in ‘Trump and Barr_e should resign’). The contextual features needed for clumping very often obtain, so clumped judgments very often follow from their unclumped counterparts.

While I cannot here provide a thoroughgoing theory of the contextual features necessary or sufficient for atomic judgments to clump, I proffer two preliminary observations. First, if ‘B does not φ ’ grounds ‘A should φ ,’ ‘A should φ ’ and ‘B should φ ’ will fail to clump. For if A should φ *because* B does not φ , then whenever B does φ , it is false that A should φ . So ‘A and B_j should φ ’ is false: φ -ing is not something which should be done jointly by A and B. Second, for two atomic judgments to clump, it is necessary but not sufficient that the acts which they prescribe be corealisable. If it is impossible for all of a set of agents to perform some act which each individually should perform, then it is false that they jointly should perform it, and clumping fails. But this necessary corealisability is insufficient for clumping because sometimes it is good for each of two agents to do a thing, and *possible but worse* for them jointly to do it. In these cases, clumping fails, not because it is impossible for A and B jointly to φ , but because it is normatively undesirable that they jointly φ . Consider this modification of our case:

LONELY AUNT*: Ann’s brother Bob has promised to visit their lonely great-aunt at the nursing home this weekend. But he is selfish and unreliable, and Ann knows he’s booked a luxury getaway. The nursing home does not place restrictions on how many visitors their great-aunt can have. But she is mentally deteriorating, and finds it overwhelming to keep up with more than one family-member. She fares best when only one relative visits.

As with the original version of the case, we judge that Bob should visit his great-aunt (because he promised to), and we judge that Ann should visit her (because Bob will fail to), but we do not judge that Ann and Bob jointly should visit her. So, ‘Ann and Bob_e should visit,’ but not ‘Ann and Bob_j

should visit.’ Perhaps a further contextual requirement, then, is that the outcome of φ -ing by all of the agents implicated in the conjunction of atomic judgments, be no worse than the outcome of one of them φ -ing.

Leaving the task of identifying all of the contextual features required for clumping for another day, it suffices here to observe that there is a semantic distinction between unclumped judgments and their clumped counterparts, such that unclumped judgments do not entail their clumped counterparts. Establishing this is enough to cause trouble for Gibbard’s analysis of third-personal judgment. The basic problem for Gibbard’s plan-expressivism, as I turn now to explaining, is that it is committed to the thesis that, whenever there is an irreducible “should” judgment, there is some planning subject with whom that judgment is associated. But clumped judgments undermine that thesis.

4.2 *The Problem for Gibbard*

A consequence of the fact that sets of normative judgments do not always clump, is that the judgment ‘A and B should φ ’ requires a two-pronged analysis from the plan-expressivist. A theory of normative judgment should tell us what attitude the unclumped judgment ‘A and B_e should φ ’ expresses, as well as what different attitude the clumped judgment, ‘A and B_i should φ ’ expresses, as well as what logical and semantic relations hold between the two. The analysis of unclumped judgments does not pose any special puzzle for Gibbard’s view. These judgments are simply the conjunction of atomic judgments, like ‘A should φ ’ and ‘B should φ .’ So their meaning is given by the standard clause for “and” in the semantic theory: for Gibbard, by taking the intersection of the semantic values of ‘A should φ ’ and ‘B should φ ’ (*viz.* ‘I plan to φ if I am A’ and ‘I plan to φ if I am B’).

By contrast, the clumped judgment ‘A and B_i should φ ’ is *not* just another way of stating the conjunction of the discrete judgments ‘A should φ ’ and ‘B should φ ,’ and it is not clear how Gibbard should understand it. For Gibbard, a normative judgment is always a contingency plan for the judging

subject to do this or that *if she is planning subject X* in circumstances *c*. The trouble posed by clumped judgments is that there seems to be no *X*. It is one thing to plan for the case in which one is Ann, or for the case in which one is Bob. But what is the contingency one is planning for when one judges that Ann and Bob *jointly* should φ ? Gibbard does not explain how clumped judgments are to be analyzed. The natural development of his view, though, would be to treat clumped judgments like any old third-personal judgment, understanding ‘A and B_i should φ ’ as a plan associated with a *conjoined* planning subject, (A and B), such that it means ‘I plan to φ if I am (A and B).’

But what does it mean to plan for being two agents at once? When I judge what another person should do, I am making a plan for what to do if I am that person. I am identifying the “thing that is to be done” from an imagined perspective: my perspective if I were the person about whom I am making the judgment. But in the class of cases we are considering, *there is no such perspective*. There is no agent whose shoes I can imaginatively fill. There is nothing that it is like to be (A and B). And there is nothing that is to be done from that practical position, because it does not, in any familiar sense, exist.⁴⁰

It may be countered that, even if there is nothing that it is like to be (A and B), there is a practical position associated with being (A and B), *viz.*, the position of a collective agent which performs various collective actions. Arguably, ‘A and B should φ ’ sometimes expresses a judgment about what a collective should do. Speaking of a doubles tennis team, the statement “Jane and Jill should regain focus so they can win the match” might mean that Jane and Jill, *as a team*, should regain their *collective* focus – where this means something different from the conjunction of judgments that Jane should regain her focus and that Jill should regain hers. So perhaps some clumped judgments can be glossed as plans associated with a collective agent. But even if this is so, the analysis will hardly

⁴⁰ Nor does it seem to be epistemically available. Michael Ridge (2018) raises this point regarding judgments about collective agents (n. 41).

apply to *all* judgments involving multiple agents, because in many cases, it is simply false that clumped judgments are saying something about a collective agent.⁴¹ For example, regarding LONELY AUNT and LONELY AUNT*, ‘it is false that Ann and Bob_i should visit’ is not the claim that it is false that Ann and Bob, understood as a collective agent, should visit. It is the claim that it is not the case that visiting should be done, independently and individually, by the both of them.

This is a subtle distinction, but a simple example elucidates it. The judgment that John and Jim ought_i to pass the exam is *not* the judgment that they ought to pass the exam *collectively* –that they ought to divvy up the labor of studying, each tackle half the questions, and so on. It is just the judgment that they both should pass. The normative considerations which support the clumped judgment ‘John and Jim_i should pass’ support the formation of two discrete intentions (John’s intention to pass, Jim’s intention to pass), and two discrete actions (John’s passing, Jim’s passing), not the formation of a shared intention or the performance of a collective action. *This* is the sense in which our clumped judgment, that they jointly should pass, is not associated with any real practical position. And this is why I deny that plan-expressivists can, as a general rule, analyze clumped judgments as contingency plans for the situation in which one is the collective agent (A and B). I leave the task of finding some superior analysis of clumped judgments – some plausible planning subject with whom to associate this class of plans – to the exponents of Gibbard’s plan-expressivism.

Gibbard’s view is not alone in its insensitivity to the clumped versus unclumped readings of third-personal judgments with conjunctive subjects, nor in failing to provide an adequate account of

⁴¹ Moreover, even those clumped judgments which can plausibly be glossed as judgments about a collective agent look to pose a problem for the plan-expressivist. Ridge (2018) rejects Gibbard’s plan-expressivism on the grounds that it cannot make sense of normative judgments about collectives. He appeals to the fact that “there is no ‘what it is like’” associated with collective agents, and therefore no “epistemically possible scenario in which [one] find[s] oneself in circumstances just like those facing” them, in arguing that it is impossible for individuals to “form a contingency plan for what to do if in the circumstances facing a collective agent” (2018: 840-844). Ridge is concerned with collective agents like states and corporations, which do not have a phenomenology, whereas the ‘collective’ we are now considering is made up of agents each of whom does have a phenomenology. But the distinction may be unimportant: there is no phenomenology associated with *being A and B simultaneously*!

the logic and semantics which govern how the two relate. Take standard deontic logic. It treats “it ought to be the case that” as a necessity operator. Necessity operators agglomerate. So standard deontic logic has it that, if it ought to be the case that $A \varphi$ and if it ought to be the case that $B \varphi$, then it ought to be the case that A and $B \varphi$. But that entailment holds for only one of the two senses of ‘ A and B should φ ,’ *i.e.*, ‘ A and B_e should φ ’. A great many theories – including a great many cognitivist theories – will have to fine-tune the semantics of “ought,” in order to theorize clumping correctly.⁴²

At the same time, plan-expressivists *are* alone in having recourse only to the menu of plan-like non-cognitive attitudes in performing this fine-tuning. For expressivists, the semantics of normative language is a function only of the non-cognitive attitude(s) in whose terms normative judgment is analyzed. This means that, when plan-expressivists wish to give an analysis of some semantic structure in normative discourse, they must locate that structure within the relevant planning attitude(s). But the semantics of clumping is subtle, so doing so could prove difficult.

To sum up: to allow for the distinction between clumped judgments and the conjunction of atomic judgments, plan-expressivists must find some ground for differentiating their analysis of judgments of the form ‘ A and B_e should φ ’ from their analysis of judgments of the form ‘ A and B_i should φ .’⁴³ This requires identifying some planning subject with which clumped judgments can plausibly be associated.

5 Conclusion

⁴² Mellor & Shea (ms.)

⁴³ Although they face other grave objections, two rivals of Gibbard’s plan-expressivism look comparatively well-positioned to deal with the semantic distinction between clumped judgments and the conjunction of atomic judgments. The first is Ayars’ version of plan-expressivism. Arguably, making decisions on behalf of a set of agents, A , B , and C , is different from making decisions on behalf of A , and separately, on behalf of B , and separately, on behalf of C . The second is preference-expressivism, which analyzes ‘ought’ in terms of a desire-like attitude. Desire-like attitudes are sensitive to the relevant distinction. One can desire to have coffee and desire to have oysters, without desiring to have coffee and oysters together (thanks to Jamie Dreier for raising this point).

Non-corealizability challenges plan-expressivism in two ways. First, plan-expressivism associates normative judgment with an attitude which satisfies a corealizability constraint, which constraint is (in the interpersonal case) unwarranted, or (in the intrapersonal case) warranted only by a contentious normative premise. On Ayars' view, interpersonal normative judgments are coherent only if they express interpersonally corealizable decisions, but this is false (§2). On both Ayars and Gibbard's views, intrapersonal normative judgments are coherent only if they express intrapersonally corealizable plans, but this is true *only on the assumption* that possibilism is true (§3). Notably, plan-expressivists have favored intention or intention-like attitudes for the analysis of normative judgment precisely because these attitudes are subject to a corealizability constraint, and therefore hold promise for the derivation of various commonsensical consistency norms on our judgments.⁴⁴ A challenge to the corealizability constraint is thus a challenge to the very feature of plan-expressivism which looks to recommend it.

Second, Gibbard's plan-expressivism associates all normative judgments with a planning subject. But cases involving non-corealizability suggest that this is a mistake, because these cases give rise to judgments about what multiple agents should do, which are not associated with any such subject (§4). In cases involving interpersonal non-corealizability, the judgments that each of a set of agents should perform some act φ do not entail the judgment that all of these agents should φ . This raises the problem of what and whose plan-like attitude this latter judgment expresses.

The arguments of this paper thus support the conclusion that normative judgments are not plan-like states. But they also call up a number of broader, contentious meta-ethical and normative questions. I shall close by identifying three of the most interesting. The first is whether a plausible

⁴⁴ Schroeder (2010). Perhaps this suggests some affinity between Gibbard's plan-expressivism and possibilism. Still, that the most prominent expressivist view rules out actualism is surprising and noteworthy. Moreover, it may be difficult to say what the commonsense consistency norms on normative judgment *are* without choosing between actualism and possibilism (§3).

account of practical normativity, and a plausible semantics for ‘ought,’ should allow for *transparently motivationally inert* normative judgments: judgments made in the knowledge that the subject of the judgment will not take them up in her intentions or actions. I have argued that normative judgments can hold in descriptive circumstances which include the fact that they will not be so taken up. For example, Claire in SICK UNCLE should visit her uncle in the hospital, notwithstanding that she will fail to do so due to her weak will. And Bob in LONELY AUNT should visit his great-aunt, notwithstanding that he will fail to do so due to his selfishness. Such judgments play a central role in everyday, rational determinations of what to do. It is partly *because* Claire should-but-will-not see her uncle, that she should stay home and phone her uncle; and it is partly *because* Bob should-but-will-not visit his great-aunt, that Ann should visit her lonely aunt.

The second is the relationship between agents’ reasons for action and their reasons for intention. In discussing intrapersonal non-corealizability, my focus has been on plan-expressivism, which *identifies* normative judgment with an intention-like attitude and – if ENKRATIC HARMONY is assumed – thereby precludes the possibility of judging coherently that an agent should φ and should ψ , where these are taken to be intrapersonally non-corealizable. Notably, there are popular views within normative ethics which arrive, *via* similar considerations, at a corealizability constraint on intrapersonal normative judgment. If ENKRATIC HARMONY is assumed, then both the view that an agent’s reasons for action *just are* her reasons for intention,⁴⁵ and the view that one is rationally required to intend in accord with her conclusive deliberative reasons, generate the same corealizability constraint.⁴⁶ This suggests that there may be interesting connections between the motivations for plan-expressivism and the motivations for such views.

⁴⁵ McHugh & Way (2022), Portmore (2011), Smith (2013).

⁴⁶ *E.g.*, as discussed above, Kieseewetter (2015) suggests that it cannot be true of an agent that she ought to perform non-corealizable acts φ and should ψ , partly *on the grounds* that she cannot rationally to intend to φ and ψ .

The third is the relationship between meta-ethics and normative ethics. It is a striking implication of my argument that the viability of plan-expressivism depends on how a controversial normative question is answered. Familiar versions of plan-expressivism presuppose and entail that it can never be the case that an agent should φ and ψ , where φ and ψ are intrapersonally non-corealizable. But this rules out a first-order view: actualism. With this point in mind, the proposal that meta-ethics and normative are isolated theoretical ambits, each to be pursued independent of the other, looks dubious.⁴⁷

⁴⁷ I am grateful to Sarah McGrath, Gideon Rosen, Joseph Shea, and Michael Smith for their extensive comments on this paper. I also thank Alisabeth Ayars, Selim Berker, Jamie Dreier, Samuel Fullhart, Mark Johnston, Rowan Mellor, Elliott Salinger, Thomas Schmidt, and Alex Worsnip; the participants of the 2022 Princeton-Humboldt Summer Institute on Practical Normativity and the 2022 Madison Metaethics Workshop; and two anonymous referees for *Oxford Studies in Metaethics*, for their helpful comments.

REFERENCES

- Ayars, Alisabeth. 2022. "Deciding for Others: An Expressivist Theory of Normative Judgment." *Philosophy and Phenomenological Research*, October, phpr.12800. <https://doi.org/10.1111/phpr.12800>.
- . "An Explanation of the Essential Publicity of Practical Reasons." *Oxford Studies in Metaethics* 19.
- Estlund, David M. *Utopophobia: On the Limits (If Any) of Political Philosophy*. Princeton, New Jersey: Princeton University Press, 2020.
- Gibbard, Allan. 1990. *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Cambridge, Mass: Harvard University Press.
- . 2003. *Thinking How to Live*. Cambridge, Mass: Harvard University Press. <https://www.jstor.org/stable/10.2307/j.ctv1pncqb8>.
- Jackson, Frank, and Robert Pargetter. 1986. "Oughts, Options, and Actualism." *The Philosophical Review* 95 (2): 233. <https://doi.org/10.2307/2185591>.
- Kiesewetter, Benjamin. 2015. "Instrumental Normativity: In Defense of the Transmission Principle." *Ethics* 125 (4): 921–46. <https://doi.org/10.1086/680911>.
- . 2018. "Contrary-to-Duty Scenarios, Deontic Dilemmas, and Transmission Principles." *Ethics* 129 (1): 98–115. <https://doi.org/10.1086/698734>.
- Kratzer, Angelika. 2012. "The Notional Category of Modality." In *Modals and Conditionals: New and Revised Perspectives*, edited by Angelika Kratzer. Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780199234684.003.0002>.
- Mellor, Rowan & Shea, Margaret. "What Are We to Do? The Ethics and Semantics of the Joint 'Ought.'" Unpublished manuscript.
- Portmore, Douglas W. *Commonsense Consequentialism: Wherein Morality Meets Rationality*. Oxford Moral Theory. Oxford, New York: Oxford University Press, 2014.
- Ridge, Michael. 2018. "Expressivism and Collectives." *Mind* 127 (507): 833–61. <https://doi.org/10.1093/mind/fzw078>.
- Roojen, Mark van. 1996. "Expressivism and Irrationality." *The Philosophical Review* 105 (3): 311. <https://doi.org/10.2307/2185703>.
- . 2018. "Moral Cognitivism vs. Non-Cognitivism." In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Fall 2018. Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/fall2018/entries/moral-cognitivism/>.
- Schroeder, Mark. 2010. *Being For: Evaluating the Semantic Program of Expressivism*. Illustrated edition. Oxford ; New York: Oxford University Press.

Smith, Michael. “3 The Ideal of Orthonomous Action, or the How and Why of Buck-Passing.” In *Thinking About Reasons: Themes from the Philosophy of Jonathan Dancy*, edited by David Bakhurst, Brad Hooker, and Margaret Olivia Little, 0. Oxford University Press, 2013.
<https://doi.org/10.1093/acprof:oso/9780199604678.003.0004>.

Timmerman, Travis, and Yishai Cohen. 2016. “Moral Obligations: Actualist, Possibilist, or Hybridist?” *Australasian Journal of Philosophy* 94 (4): 672–86.
<https://doi.org/10.1080/00048402.2016.1140789>.

White, Stephen J. 2017. “Transmission Failures.” *Ethics* 127 (3): 719–32.
<https://doi.org/10.1086/690014>.

Way, Jonathan, and Conor McHugh. “All Reasons Are Fundamentally for Attitudes.” *Journal of Ethics and Social Philosophy* 21, no. 2 (February 28, 2022).
<https://doi.org/10.26556/jesp.v21i2.1341>.

Worsnip, Alex. 2019. “‘Ought’-Contextualism beyond the Parochial.” *Philosophical Studies* 176 (11): 3099–3119. <https://doi.org/10.1007/s11098-018-1165-9>.

———. 2021. *Fitting Things Together*. Oxford University Press.
<https://doi.org/10.1093/oso/9780197608142.003.0007>.