ORIGINAL ARTICLE

Deciding for Others: An Expressivist Theory of Normative Judgment

Alisabeth Ayars ⓘ

University of British Columbia

Correspondence
Alisabeth Ayars, University of British
Columbia
Email: alisabeth.ayars@ubc.ca

Abstract

This paper develops a new form of metaethical expressivism according to which the normative judgment that X should Φ consists in a decision that X Φ. When the judgment is first-personal—e.g., my judgment that *I* should Φ—the view is similar to Gibbard's plan expressivism, though the state I call "decision" differs somewhat from a Gibbard-style plan. The deep difference between the views shows in the account of third-personal judgments. Gibbard construes the judgment that Mary should Φ as a de se plan on the thinker's part to Φ if she turns out to be Mary (the Subtle View). I construe the judgment as a decision for Mary *that Mary Φ* (the Simple View). The main argument for Simple Plan Expressivism is that it solves problems for Gibbard's approach, resonates with a new and interesting moral psychology, and better makes sense of certain independently plausible constraints on normative judgment. In the end I argue that this account of normative judgment has implications for first-order ethics, implying in particular that rational egoism as standardly formulated is incoherent.

1 | INTRODUCTION

In their effort to provide a systematic account of normative thought and discourse, expressivists have for the most part taken first-personal normative judgments as their primary analysanda. Suppose after my husband cheats on me, I conclude: "I should leave my marriage." According to

---

expressivists, my first-personal judgment that I should leave my marriage consists in a desire-like attitude toward leaving my marriage: e.g., a preference (Dreier 1996, Silk 2015[1]), an intention, or a *sui generis* conative state. For instance, on Allan Gibbard's influential view (2003), my judgment consists in a plan or decision to leave my marriage, a state that settles for me the practical matter of what to do.[2]

My focus in this article is not moral judgments, but judgments about what people "should" or "ought" to do *all things considered*—the sort of normative judgments that close deliberation and which we express when we give general normative advice. Expressivists aim to capture the practical, motivating character of normative judgment, and when it comes to first-personal judgments about what we ought to do, an expressivist analysis is attractive, since these judgments have an obvious action-guiding role. However, we also make third-personal normative judgments: judgments about what other people should do, all things considered. Gossiping with you, I might say, "Mary should leave her marriage." This judgment is about what Mary should do, rather than what I should do. Since expressivists must account for the full range of normative judgments, they need an account of these third-personal judgments, and that turns out to be surprisingly difficult to provide.

The simplest view the expressivist could take is that just as my first-personal judgment that I should leave is a desire-like attitude toward leaving my marriage, my third-personal judgment that Mary should leave is a desire-like attitude of the same type toward *Mary*'s leaving her marriage. For instance, if my judgment that I should leave is a preference for my leaving over not leaving, my judgment that Mary should leave is a preference for Mary's leaving over her not leaving. If my judgment that I should leave is a decision for me (a decision for me to leave), my judgment that Mary should leave is a decision for Mary: a decision that Mary leave; and so on. Call this the

> **Simple View of Third-Personal Judgments:** Just as my first-personal judgment that I should Φ is a desire-like attitude toward my Φing, my third-personal judgment that X should Φ is a desire-like attitude of the same type toward X's Φing.

The Simple View has a virtue: It is uniform, in the sense that it represents first-personal and third-personal judgments as states of exactly the same kind, differing only in the content of the underlying attitude. But the Simple View is widely rejected, in part because it appears to misdescribe the motivational role of third-personal judgments. On the Simple View the third-personal judgment is a plan, or a desire, or a preference on my part concerning Mary's conduct; but any such state would have to motivate the subject to promote its object, or so it seems. And the trouble is that third-personal normative judgments are not in general motivating in this way. If Mary is my friend, the judgment might come along with something like a desire to help her leave. But I can make the judgment even if Mary is a stranger, or an enemy, in which case it will not incline me to help her leave. This problem arises for any version of the Simple View. But for proponents of the main contemporary form of expressivism, plan expressivism, there would be a further problem. The plan expressivist says that normative judgments

---

are decisions. So, if she adopts the Simple View, she will say that my judgment that Mary should leave is a decision on my part *that Mary leave*. But there seems to be no such thing as deciding for another person. So the plan expressivist looks to have two good reasons for rejecting the Simple View.

The main alternative expressivist treatment of third-personal judgments involves identifying them with *conditional* first-personal attitudes (Gibbard 1983, 2003). My judgment that Mary should leave may not motivate me to help her leave. However, it would, plausibly, motivate me to leave my marriage *if* I were in circumstances just like Mary's in all relevant respects. This suggests the desire-like attitude expressed by a third-personal judgment is conditional and first-personal: it is a desire-like attitude towards my leaving if, hypothetically, I am in Mary's situation. On Gibbard's view, it is a conditional plan: a plan to leave should I find myself in Mary's exact predicament.

Suppose Holmes is reasoning about what Mrs. Hudson should do. Holmes cannot decide *for* Mrs. Hudson, says Gibbard. But he can form a contingency plan for what to do in her situation:

> [I]f the two aren't in communication tonight [Holmes] knows that his own thinking can in no way settle the question of what to do in Mrs. Hudson's own mind. She can't possibly act on his thinking, the way Holmes later on can act on the planning he's now engaged in. Still, this is in a way a kind of hypothetical planning: he thinks what to do if in a hypothetical circumstance. He's thinking as if he could plan what to do if in Mrs. Hudson's plight ... (Gibbard 2003: 50).

Even though Holmes knows that he will never be Mrs. Hudson, he can still make a hypothetical plan for himself for what to do in her shoes, and according to Gibbard, this is what his third-personal judgment that Mrs. Hudson should Φ consists in. Call this the

> **Subtle View of Third-Personal Judgments**: My first-personal judgment that I should Φ is desire-like attitude toward my Φing. My third-personal judgment that X should Φ is a desire-like attitude toward *my Φing in a situation exactly like X's.* The first part of this paper is critical: I argue that Subtle View of third-personal judgments is inadequate. In the second half, I defend an approach to third-personal judgments within a plan expressivist framework according to which third-personal judgments are unconditional plans or decisions. Specifically, I argue for *Simple Plan Expressivism*, according to which the third-personal judgment that X should Φ is a decision for X *that X Φ*—a decision whose focus is someone else's conduct. This account avoids all the objections to the Subtle View that I set out in the first half the paper.

This is a surprising view—the notion of deciding for another person seems unfamiliar. But, I argue, decisions for others are operative in certain widespread social practices like *delegated decision-making*, in which one person allows another to settle a practical question for her. Thus, decisions for others are not only possible, they are commonplace. Moreover, I argue that decisions for others do not motivate the psychological subject of the decision (the decider) to promote their object, but rather function well when they motivate the *executive subject* of the decision—the agent whom the decision is a decision for (Roth 2004). The account is thus consistent with the motivational asymmetry noted earlier, according to which first-personal judgments motivate the subject of the judgment to promote their object, whereas third-personal judgments in general do not. Moreover, the view has the interesting and intuitively plausible consequence that the motivational role of third-personal judgments is essentially social, aimed at guiding the actions

of others.[3] I will not pretend to offer a full-dress defense of my approach in this paper. My aim is to set out the view, to highlight its attractive features, and to show why it is superior to Gibbard's account of third-personal judgments.

Finally, I argue that this account of normative judgment has implications for first-order ethics. If my third-personal judgments are decisions I make on behalf of others, then they must harmonize with one another and with my decisions for myself in the same way in which my decisions about what to do now must harmonize with my decisions about what to do at other times. This coherence constraint on decisions restricts the normative judgments we can coherently make, implying in particular that rational egoism—the view that everyone should benefit himself—is not just false but (in a sense I specify) inconsistent. The view thus promises a route from metaethics to ethics and holds out the prospect that reflection on the nature of normative judgment may yield substantive constraints on what we should do.

2 | PROBLEMS FOR THE SUBTLE VIEW

I will illustrate the problems with the Subtle View against the background of Gibbard's plan expressivist theory since Gibbard provides the most developed version of the Subtle View. However, the objections apply to any expressivist view which identifies third-personal judgments with conditional first-personal attitudes, like conditional desires or preferences.

Before setting out my three main objections to the Subtle View, I will state a general concern. As noted, Gibbard holds that normative judgments are plans, and in particular that my third-personal judgment that Mary should leave is a plan for me to leave if I find myself in Mary's situation. Many have wondered about the psychological plausibility of the view that third-personal judgments are plans or preferences for remote scenarios (e.g., Cullity 2007, Gregory 2017). Why would we bother forming plans for situations that we know will never arise? According to Gibbard, we engage in such planning as a kind of "rehearsal for life" (Gibbard 2003: 52)—to settle in advance what we will do in various contingencies. The capacity for third-personal judgment emerged in human beings because this sort of rehearsal proved advantageous.

However, the relevant scenarios are often extremely remote; we can make normative judgments about what the president of Venezuela ought to do, or what a character in a fantasy novel ought to do. It is hard to believe that our hominid ancestors derived a practical benefit from by simulating decision-making in situations they will certainly never face. Especially noteworthy is the fact that on Gibbard's view, Holmes is not merely planning to Φ in a situation qualitatively identical to Mrs. Hudson's, but planning to Φ if he is *Mrs. Hudson* in Mrs. Hudson's shoes (Gibbard 2003: 51). If he judges she should leave, he is issuing himself the (conditional) directive: "if you're ever Mrs. Hudson in the exact situation she's in right now, leave." Since it is impossible for Holmes to be Mrs. Hudson, this plan has no chance of being relevant to a situation Holmes finds himself in—not even a qualitatively similar situation (since Holmes would still be Holmes in this situation).

---

[3]Of course, there are a variety of rich and highly developed descriptivist, truth-conditional semantics for "should" which support cognitivism about these judgments. See, for instance, Wedgwood (2006, 2007: ch. 5, 2016), Finlay (2009, 2010, 2014), and Worsnip (2019). Unfortunately, I shall not be able to address any of these cognitivist alternatives here. My aim in this paper is to investigate the best response to the problem of third-personal judgments within an expressivist framework.

While I share these authors' skepticism on this score to a degree, I think this objection is not decisive. Perhaps we have come to find such contingency planning "fascinating in itself, entirely apart from the ways the practice might later come in handy" (Gibbard 2003: 53). But psychological plausibility is not the only concern for the Subtle View; so let us turn to the other objections.

2.1 | The problem of interpersonal disagreement

My first objection is that the Subtle View does not give a satisfactory account of interpersonal normative disagreement. Suppose you and I are gossiping together about Mary in the coffee shop and I judge that Mary should leave her marriage, while you judge that it is not the case that she should leave. Our normative judgments obviously disagree, so the mental states that correspond to them must disagree. Moreover, our normative judgments disagree in the strong sense that they are inconsistent with one another. This is a datum that any account of normative judgment must capture.

On Gibbard's view, my judgment that Mary should leave is a plan that requires *me* to leave in Mary's situation, whereas your judgment that it's not the case that Mary should leave is a plan that permits *you* not to leave in that situation. Now if a single person made both these plans, his plans would be inconsistent, in the sense that one would forbid an action that the other permits. At least, this is what plan expressivists must say in order to account for why it would be inconsistent for a single agent to judge both that Mary should leave and that it's not the case that Mary should leave. But a special problem arises in the case of interpersonal disagreement. When I make the first plan and you make the second, my plan requires *me* to leave in Mary's shoes, whereas your plan says nothing about my conduct, and so neither forbids nor permits nor requires *my* leaving in Mary's shoes. The disagreement thus goes missing; my plan does not forbid any course of action that your plan permits.

Of course there is a sense in which our plans "disagree" on the Subtle View. We make de se plans that would be inconsistent if a single person made them. But that's analogous to the sense in which we "disagree" if I judge that I'm happy and you judge that you are not. These judgments "disagree" in the sense that they respectively self-ascribe properties that no single person could have. But of course they don't really disagree. If I say "I'm happy" and you say "I'm not", though we are accepting and rejecting (respectively) the same centered proposition, we do not contradict one another in any sense. Gibbard's view entails that when I say "Mary should leave" and you say "No, she shouldn't", we disagree only in this bogus sense. Conflicts in de se contents do not amount to disagreement.[4]

The worry that I am raising is not the familiar "negation problem" for expressivism (Unwin 2001, Schroeder 2008).The negation problem is the task of explaining why it is inconsistent even for a single individual to judge both that X should Φ and that it's not the case that X should Φ. Gibbard has a solution to the negation problem. As both Gibbard and I understand them, plans are states that reject or permit actions. A plan requires an action if it permits only actions consistent with it; a plan forbids an action if it permits only actions inconsistent with it. When an agent judges that X should Φ, she accepts a plan that requires her to Φ in X's shoes; when she judges

---

[4]See Dreier (2009) for a further discussion of expressivists' problems with disagreement, including a discussion of Gibbard. For further discussion of why, in general, conflicts of de se contents do not suffice for disagreement, see Beddor (2018); Egan (2010, 2014); Richard (2011), and Ross and Schroeder (2013). See Ridge (2013) and Beddor (2018) for accounts of disagreement which might be marshalled in defense of Gibbard.

that it's not the case that X should Φ, she accepts a plan that permits her not to Φ in X's shoes. And it can be shown, given plausible definitions, that these plans are inconsistent.[5]

The problem I am pressing is that when you and I disagree about whether Mary should leave, our Gibbard-style plans are not inconsistent in the same way: Yours requires *you* to leave in Mary's shoes; mine permits *me* not to leave. These are *different* plans, but this sort of difference does not amount to disagreement. If I plan to go to the beach today and you plan either to go to the beach or stay home—embracing a plan that permits either—we have different plans, but there is nothing we disagree about since your plan concerns how you are to spend the day, and mine says nothing about that.

2.2 | The wrong subject matter objection

My second objection to the Subtle View is that it assigns the wrong subject matter to third-personal normative judgments.

When I'm gossiping with you about whether Mary should leave her marriage, we are talking *about* Mary and *only* Mary. My judgment that Mary should leave is focused on Mary in the same sense in which my ordinary descriptive belief that Mary is happy has Mary as its sole focus. The judgment is not at all about me, just as my belief that Mary is happy is not at all about me.

The Subtle View entails that my third-personal judgment is a de se attitude, and is therefore partly about myself. I say "partly," because the attitude does make reference to Mary. As Ridge points out,

> [I]ndexical attitudes can take other agents within their scope in the following way: I can prefer that if I were in the other agent's circumstances that certain states of affairs not obtain. The attitude is still indexical in that it is cast in terms of what I want for myself if I were in other person's circumstances. At the same time, it makes a kind of sense of the idea that the judgment is nonetheless in another sense about the other agent in question (Ridge 2011: 48)

However the Subtle View cannot vindicate the idea that third-personal judgments are *solely* about the agent and not the subject of the judgment. Consider the difference between my belief that Zoe is nervous and my belief that if I were Zoe, I would be nervous. While the former is solely about Zoe, the latter is partly about me. The Subtle View thus makes judgments about others partly self-regarding in just this way, and that is false to the phenomenology of normative judgment.

---

[5]A Gibbard-style plan is a set of hyperplans, where a hyperplan specifies, for each centered world, one or more actions permitted to the center. A plan permits X to Φ if every element of it permits the center to Φ if the center is X; a plan requires X to Φ iff it does not permit the center not to Φ if the center is X. Plans are inconsistent when their intersection is empty. Since there is no hyperplan that both requires the center to Φ if the center is X and permits the center not to Φ if the center is X, the plans associated with "X should Φ" and "It's not the case that X should Φ" are inconsistent and any single person who accepts both is thereby inconsistent. But note, since these plans are de se contents, built from centered worlds, it will not follow from the fact that you accept one plan, and I accept an inconsistent plan, that we disagree; conflicts in de se contents do not amount to disagreement. That is why the disagreement problem is different from the negation problem. As many have pointed out, Gibbard's formal solution to the negation problem must be supplemented with an independent account of what it is for one state of mind to disagree with another (Schroeder 2008, Dreier 2009). My point is that even if we grant that such an account can be provided, it would not solve the *interpersonal* disagreement problem.

One might reply that actually Gibbard gets it right—that third-personal normative judgments *are* in fact about oneself—by noting that when we give advice, we often say things like "if I were you, I wouldn't do that." (Indeed, Gibbard relies heavily on this point).[6] However, this is weak evidence for the claim that the judgment expressed by "you should Φ" is *equivalent to* the judgment expressed by "if I were you, I would Φ". The fact that I think about what I would do if I were you in giving advice is easily explained in other ways. For instance, it is plausibly explained by the fact that it is often easiest to figure out what someone should do in a situation by thinking about what we would do. This is true of factual questions as well. If I want to figure out whether Zoe is nervous, I might do so by thinking about whether I would be nervous in Zoe's situation. In a context in which I'm asked whether I think Zoe is nervous, I might say, "Well, if I were Zoe, I'd be nervous." But this is hardly evidence for the claim that the judgment that Zoe is nervous is *identical* to the judgment that if I were Zoe, I would be nervous.[7]

2.3 | The conflation objection

The third objection is that the Subtle View conflates normative judgments that are distinct, in the sense that it is possible to hold one without holding the other (Gregory 2017[8]). Consider the following two judgments: My judgment that Mary should Φ and my judgment that I should Φ if I am Mary in Mary's circumstances. The Subtle View presumably gives these judgments the same analysis: they are both desire-like attitudes toward Φing if one is Mary in Mary's circumstances. In other words, according to the Subtle View, third-personal normative judgments are first-personal normative judgments in disguise. But these judgments are distinct: I can accept one without accepting the other. Gregory does not give an example of a scenario in which someone accepts one but not the other, but we can construct one. Suppose I'm a selfish egomaniac who thinks the following: "In every possible situation, everyone should do what is best for me." I then accept the judgment that if I am Mary in Mary's circumstances, I should do what is best for Mary.[9] However, I reject the unconditional judgment that Mary should do what is best for her-

---

[6] See, e.g., Gibbard 2002.

[7] The theoretical arguments of Lewis (1979) may show that every seemingly *de dicto* judgment is in fact implicitly *de se.* But a version of the objection would still apply. Even in Lewis's framework we can distinguish the *de dicto* judgment that Mary should Φ — which amounts to self-ascribing the property of being such that Mary should Φ, a property that never differs among world-mates— from the properly *de se* judgment that I should Φ if I am Mary, which amounts to self-ascribing *x. x should Φ if x is Mary in Mary's shoes*, a property that can differ among world-mates. The objection would then be that Gibbard's proposal represents third-person normative judgments as *doubly de se* — as self-prescribing a plan *for oneself* — when intuitively the content of the judgment involves no anaphoric back reference to the subject of the attitude.

[8] Gregory presents a version of this objection (Gregory 2017: 43-44). Gregory's version uses the subjunctive conditional, rather than the indicative conditional. I am using the indicative conditional because it is clearer that indicative conditionals are analyzed as contingency plans in a Gibbardian framework. Indeed, Gibbard often states contingency plans in indicative form; for instance, he says that Brutus, in planning to stay home in Caesar's shoes, thinks "If I am Caesar in Caesar's shoes this morning, let me stay home" (Gibbard 2003: 6).

[9] The example may be more compelling if we imagine that the egomaniac does not know whether she is Mary. The indicative is infelicitous when the subject knows the antecedent is false. If the egomaniac does not know whether she is Mary, then she accepts the judgment expressed by "if I am Mary in Mary's circumstances, I should do what is best for me" but does not accept the judgment that Mary ought to do what is best for herself, since she does not know whether she is Mary. For more discussion of the conflation point and the egomaniac, see (Ayars & Rosen, forthcoming).

self, since I know that I am not Mary. Since I'm an egomaniac, I think Mary should be doing what's best for *me*.[10]

A central desideratum for any non-cognitivist theory is that it avoids conflating distinct normative judgments. If P and Q are normative claims that can be accepted and rejected independently of one another, then the theory should assign P and Q acceptance states which can be held independently of one another.[11] The Subtle View of third-personal judgments flouts this desideratum.[12]

3 | THE SIMPLE VIEW SOLVES THE PROBLEMS FOR THE SUBTLE VIEW

The Simple View, in contrast, is not subject to these objections. This is because the Simple View does not represent third-personal judgments as implicitly de se but rather identifies them with attitudes held directly toward others' actions.

Begin with the *wrong subject matter objection*. Recall that according to the Simple View, my first-personal judgment that I should Φ is

(D1) a desire-like attitude toward my Φing

whereas the third-personal judgment that Mary should Φ is

(D2) a desire-like attitude toward Mary's Φing

(In contrast to a desire-like attitude toward Φing in Mary's shoes, as the Subtle View posits). This is a unified view according to which the attitude expressed by "X should Φ" is about X and only X in both the first-personal and third-personal case. Hence, the wrong subject-matter objection is avoided.

The *problem of interpersonal disagreement* is also solved. Someone who thinks Mary should leave disagrees with someone who thinks that it is not the case that Mary should leave in virtue of their having inconsistent desire-like attitudes toward *Mary's* leaving. If the relevant attitudes are plans, the judgment that Mary should leave is a plan *that Mary leave* (i.e., that forbids Mary's not leaving). Someone who thinks Mary should leave has a plan that forbids Mary's not leaving, and someone

---

[10]Some philosophers think Qualitative Supervenience—the claim that there are no two possible situations alike in all non-normative qualitative respects but which differ in some normative respect—is a conceptual truth. If they're right, accepting one judgment but not the other is a conceptual confusion. But what matters is whether accepting one judgment and not the other, is *possible*, not whether it is coherent; see Gregory (2017) for discussion. In addition, Gibbard denies Qualitative Supervenience, since he believes one can plan differently for being (say) Jane and for being Mary (Gibbard 2003: 50).

[11]The acceptance state for a statement S is the mental state in which the acceptance of S consists, i.e., roughly, the state that would normally be expressed by a sincere assertoric utterance of S (Dreier 1996).

[12]I concede that this test of the distinctness of the judgments (involving the egomaniac) is not decisive. It certainly seems possible to judge that Hesperus is rising without judging that Phosphorus is rising; but sophisticated solutions to Frege's puzzle sometimes say otherwise (Soames 1989). Still it counts against a theory if it identifies judgments that seem on informed reflection to be modally separable.

who thinks that it's not the case that Mary should leave has a plan that permits Mary's not leaving. The plans thus disagree in a straightforward sense. On the Simple View, disagreement between individuals is like disagreement within a single individual. Insofar as expressivists can give a satisfactory account of intrapersonal normative disagreement, interpersonal disagreement comes for free. This is a major advantage of the Simple View.

Finally, the view solves the *conflation objection*: it distinguishes the judgment that Mary should Φ from the judgment expressed by "I should Φ if I am Mary in Mary's situation," identifying the former with a desire-like attitude toward Mary Φing (e.g., a plan that Mary Φ) and the latter with a desire-like attitude toward Φing if one is Mary in Mary's shoes (e.g., a plan to Φ if one is Mary in Mary's situation). The egomaniac has a desire-like attitude toward benefiting Mary if he is Mary, without having an unconditional desire-like attitude toward Mary's benefiting Mary. The view thus assigns distinct acceptance-states to these distinct normative judgments.

The Simple View immediately solves the objections to the Subtle View. So it would be nice if the Simple View were tenable. But is it?

4 | SIMPLE PLAN EXPRESSIVISM

Whether the Simple View is tenable, I believe, depends on *which* desire-like attitude the expressivist takes normative judgments to consist in. Any view that identifies third-personal judgments with preferences or desires that others do things is a nonstarter; it is just obvious that you can judge that someone should do something without preferring or desiring that she do it. Consequently, in the remainder of the paper, I will defend a view of third-personal normative judgments within a plan expressivist framework according to which third-personal judgments are unconditional *decisions* for others that others do things, as opposed to conditional decisions for the subject to do things in remote situations, as the Subtle View posits. Call this

> **Simple Plan Expressivism:** Just as my first-personal judgment that I should Φ is a decision that I Φ, my third-personal judgment that X should Φ is a decision that X Φ.

Simple Plan Expressivism gives third-personal judgments a uniform analysis, since a first-personal decision to Φ is a special case of a decision-*that*—a decision that one Φ.[13] It is hence a version of the Simple View and avoids all the objections to Gibbard's Subtle Plan-Expressivist View of third-personal judgments.

I will say why I think Simple Plan Expressivism is promising shortly. But first let me make a preliminary clarification. It is important for the plan expressivist to insist on a distinction between the state we are calling "decision" or "plan" and an ordinary intention. (This involves a departure from ordinary language, where to decide to Φ just is to form the intention to Φ). As many critics of Gibbard have noted, some intentions are obviously not normative judgments. We sometimes *pick* between options which we believe to be equidesirable. Facing a rack of practically identical black dresses that I might wear to the cocktail party, I pick the one directly in front of me. But there is no judgment that I *should* pick that one over the others on the rack, since I believe them all to be equidesirable. Call this a *Buridan case*. Sometimes we form intentions akratically, as when I decide to get drunk at the cocktail party while knowing I should not.

---

[13]English grammar omits the subject in infinitive clauses, so we speak simply of my decision to Φ. But it is clear that every such planning state consists in a system of decisions for the planner to do this or that in such and such conditions.

Assuming I am not of "divided mind,"[14] the akratic intention is not a normative judgment. Bratman calls this the *restriction problem* for plan expressivism: "Gibbard's theory needs a way of restricting which decisions or contingency plans are normatively expressible. It needs to explain why specific decisions (plans) [like the plans in Buridan cases] are not expressible normatively, even though the state of mind with which his expressivism starts is that of decisions (plans) to act in certain circumstances" (Bratman 2006: 708).

In light of the restriction problem, the plan expressivist should not identify normative judgments with intentions at all. Rather, she should identify them with a *sui generis* decisional state whose characterization must proceed, I believe, in functionalist terms. *A decision is the state that concludes practical deliberation.* The essential role of decision (and hence normative judgment) in our psychological economy is not to represent an independent normative fact, as the cognitivist would have it, but to settle or constrain our intentions up to arbitrary picking in a fashion responsive to our judgments about our reasons.[15] We might think of a decision as a state that is posterior to deliberation and just prior to intention: a decision closes practical deliberation and then is *executed* by forming an ordinary intention. Sometimes we give ourselves a menu of options to pick from, as when, at the clothing store, I think: "any one of these dresses would be okay." The critical point is that the well-functioning of normative judgments consists (at least in part) in their issuing in an intention in line with the selection, and hence their direction of fit is (at least in part) world-to-mind. This is the sense in which normative judgment is a practical state deserving of the name "decision".[16,17]

Accordingly, Simple Plan Expressivism is not the view that third-personal judgments consist in *intentions* that others do things. The view is that third-personal judgments consist in a *sui generis* state with a certain functional role, which—for reasons that will become apparent— deserves to be called a decision for another.

To know whether Simple Plan Expressivism is tenable, we need to address the objections stated earlier. The first objection was that the very idea of deciding for another person makes no sense. Decisions answer the practical question of what *to* do, and function well when they issue

---

[14]Gibbard argues that akrasia does involve divided mind. If you judge that you ought to Φ but fail to intend to Φ, "there's a part of you that doesn't really think you ought to. You are of more than one mind on whether you [ought to Φ]" (2003: 153). The solution I am offering departs from Gibbard's treatment of akrasia. My solution vindicates the possibility of weakness of will as traditionally understood: intentional action that is contrary to one's (wholehearted) judgment about what one ought to do.

[15]What are judgments about reasons in a plan expressivist framework? Gibbard understands the notion of a reason in terms of the psychological notion of "weighing a factor in favor" of an act. To judge that R is a reason to do A is to "form contingency plans on a certain pattern," a pattern which involves weighing R in favor of doing A in deliberation about what to do (2003: 190). See Bratman (2000: 257-258) for view according to which judgments about reasons consist in deliberative policies.

[16]See, e.g., Dreier (2015), who says that "it is part of the essential role of a normative judgment that it motivates" (164). Normative judgments are *supposed* to motivate in some sense (2015: 164); this is what distinguishes their role from that of cognitive states.

[17]Here is an analogy. Suppose two people are on a road trip in an unfamiliar place. The passenger, map in hand, is in charge of navigation. Her task is to instruct the driver where to go. She surveys the reasons for and against taking various routes and deploys these judgments to arrive at instructions to the driver: "turn left here," "take the next exit," "going either right or left would be okay," etc. The driver's role is to execute the passenger's instructions, by setting the car in a certain direction. When all goes well, the driver dutifully follows the passenger's instructions. There is a reasonable sense in which the passenger is deciding where the car goes, even though she doesn't have control of the wheel. The passenger settles the *question* of what route to take, in light of the reasons, by ratifying one or more options; her instructions "function well" when they are appropriately executed by the driver. In this scenario, the passenger is playing the role of normative judgment, and the driver is playing the role of intention.

in an intention and thus motivate the agent to act. In the first-personal case, my decision ratifies one or more courses of action. When all goes well, I form an intention by picking one of them. The worry is that nothing comparable is possible in the third-personal case. How could I be in a state whose role is to resolve a practical question for *someone else*? My mental states cannot reach into the brain of someone else and cause an intention to materialize. It seems that there is an *own-action condition* on decisions: one can only decide one's own actions. Call this the

> *First Objection to Simple Plan Expressivism:* Other people's actions are not under our control; our mental states do not issue in intentions on the part of others. But a decision is a state whose role is to issue in an intention. So, there is no such thing deciding for another.

The second objection was that Simple Plan Expressivism appears to misdescribe the motivational role of third-personal judgments:

> *Second Objection to Simple Plan Expressivism*: We can judge that someone should do something without being motivated to promote their doing it. But Simple Plan Expressivism arguably implies that anyone who judges that X should Φ is motivated to promote X's Φing, since decisions seem to be states which motivate the subject to promote their object.

In the next section, I will give a more thorough exposition of Simple Plan Expressivism. Then I will show, contra the First Objection, that we can and do settle practical questions for other people, and that contra the Second Objection, Simple Plan Expressivism does not entail that someone who judges that X should Φ is thereby motivated to promote X's Φing.

5 | THE OBJECTIONS TO SIMPLE PLAN EXPRESSIVISM RECONSIDERED

Simple Plan Expressivism says that third-personal normative judgments consist in decisions for others, rather than contingency plans to do things in others' shoes. But can we really make decisions for others? And even if we can, why think third-personal normative judgments consist in such decisions? I will begin by discussing an advisory social practice, which I call *delegated decision-making*, in which decisions for others are operative, and in which these decisions can be expressed using normative terms.

Suppose you and I are at a restaurant, and you know more about the menu and what I would like than I do. So, I ask you to decide what I will order: "Tell me what to get," I say. You deliberate, taking my tastes, my budget and the rest into account; and then voice your conclusion with an imperative— "Get the burrito"—or a normative assertion: "You should get the burrito". Either way, the judgment aims to settle for *me* the practical question of what to order in a fashion sensitive to my choice situation. I asked you to decide on my behalf what I am to get, and you did. If the exchange goes smoothly, I act on your decision. Your decision issues in an intention on my part to get the burrito, much as my own decisions (normative judgments) issue in intentions for me when all goes smoothly.

Of course, the process by which the decision issues in the intention involves additional steps in the third-personal case. You must communicate your decision to me, for instance. There is no magic by which your decisions automatically produce intentions in me. Still, the *role* played by your decisions is in relevant respects the same as the role that my own first-personal decisions

play in my psyche: to guide and give rise to action on my part. In delegated decision-making, instead of deliberating myself about what to do, you deliberate for me: thus it is *you* who takes my reasons into account and settles the practical question of what I am to do by providing the menu. So, it is your mental state which fulfills the relevant role.

The essential feature of delegated decision-making, as I have characterized it, is that person A delegates a decision to person B who then *decides* (on behalf of A) *that* A is to do something. B hence makes a real decision, which has as its fulfillment condition A's performing some action. Not all instances of reason-giving or persuasion involve decisions for others in this sense. If you merely point me to a consideration in favor of getting the burrito, you have provided me with a reason (by bringing it to my attention), but you have made no decision. You are not in a state whose aim is to *settle* my intentions. You are only offering input.

There are two main sources of support for the claim that delegated decision-making involves deciding for others in this sense. The first is the role played by the output of delegated decision-making, as just discussed. Your decision functions well when it motivates *me* to act, and hence is plausibly a decision for me.

Another piece of evidence that you are making real decisions that refer to my actions in this context. Notice that certain *rational norms* characteristic of decisions appear to govern your states.[18] Suppose the restaurant only has one burrito left, and that you know this. If you decide on my behalf that I am to get it, you cannot consistently also decide to get the last burrito for yourself. If after telling me I should get the last burrito you say to the waiter, "I'll get that last burrito," I would be puzzled, and could reasonably complain: "I thought you said *I* should get the burrito." This makes sense if your state is a real decision for me on your part, since if it is a real decision, it must harmonize with your other decisions, including your decisions for yourself, in the same way that your decisions about what to do now constrain your decisions about what to do at other times. If you decide to Φ, and also to Ψ, it must be possible for you to both Φ and Ψ. And more generally, if you decide (in our technical sense) that X is to Φ and that Y is to Ψ, then it must be possible for X to Φ and Y to Ψ. (This *co-realizability constraint* is the aspect of this view that makes trouble for rational egoism, as I will show later).

Notice that on the Subtle View of what your normative assertion ("you should get the burrito") expresses, there is no incoherence in making the assertion and then deciding to order the burrito yourself. Your judgment, on the Subtle View, is a conditional decision to order the burrito in my shoes. Since you are not in my shoes, you are rationally free to form the conditional plan to get it in that (unrealized) situation and then to decide to do something incompatible in the situation you actually face. This counts in favor of the present view, against the Subtle View.[19]

---

[18]Roth argues on similar grounds that we can form intentions that others do things (Roth 2004, 2016). I am, however, making a claim about the state associated with normative judgment, which is not intention in the ordinary sense.

[19]To convey the feel of how the view is different from Gibbard's, consider an example in which someone decides that she will Φ if she is in X's situation without deciding for X that X Φ. You could decide to get the burrito in my shoes while withholding judgment on what I should do. You might express this by saying, "I don't know what you should do—I'm not going to decide for you—though if *I* were you, I'd get the burrito." Your conditional decision would commit you to a motivation to get the burrito in my shoes. But it would not commit you to deciding not to interfere since it falls short of an endorsement of my getting the burrito (assuming you genuinely withhold judgment on what I should do).

However, someone who judges that if she were in X's shoes, she should Φ but refuses judge that X should Φ strikes as cagey or unreasonable. This is because normative judgments are not supposed to be tied to particular people. Any reason to favor your getting the burrito in my exact situation (which includes all of my qualitative properties) would also be a reason to favor my getting the burrito. Judging that you should get the burrito in my shoes *commits* you to judging that I should get the burrito, and this is why the assertion described sounds odd.

Contra the First Objection—the complaint that there is no such thing as settling a practical question for someone else—the example of delegated decision-making shows that we can and do make bona fide decisions that refer to others' actions, and these decisions are normatively expressible. We can intelligibly address, and sometimes resolve, the practical question of what *someone else* is to do. Some of the considerations I am invoking in support of this should be familiar. Philosophers have argued, on similar grounds, that we can form *intentions* which refer to other peoples' actions (Searle 1990, Gilbert 1996, Bratman 2013) and that the rational norms characteristic of intentions can extend across people (Roth 2004, 2014, 2016, 2018). I am, however, advancing a novel claim about the state associated with *normative judgment*. As noted earlier, for the plan expressivist normative judgment is not intention in the ordinary sense. My claim is that we can answer practical questions for others, and thus "decide for others" in the heavyweight sense of *decision* whose mark is aptness for expression in normative terms: e.g., "you should get the burrito".

Delegated decision-making is rare, and if this were the only clear context in which we can make decisions for others which fulfill their role, the point might be of limited interest. But this is not the only such context. Plausibly, the paradigmatic normative exchange is *co-deliberative*, rather than purely deferential as in delegated-decision making. Normative judgment paradigmatically involves reasoning *together* with someone about what she is to do. That is, I concur with Manne's (2014) suggestion that normative judgments are most at home in the "interpersonal" mode of interaction, i.e., within a certain kind of commonplace social practice in which one reasons *with* someone about what she is to do. If I judge that you should Φ and tell you what I think and why I think it, you may come to agree with my judgment. In agreeing with my judgment, you make a decision for yourself: a decision to Φ.[20] If all goes well, you then act on the (joint) decision. My decision for you fulfills its role indirectly by securing agreement with you.

Of course, we can make third-personal normative judgments outside these "cooperative" social practices. We can make normative judgments about strangers and enemies, with no expectation that they will be acted on. We can make normative judgments privately, without expressing them at all. Simple Plan Expressivism says that these normative judgments consist in the *same underlying state* as is present in delegated decision-making and co-deliberation—a decision for another—even if there has been no delegation of authority and even if there is no cooperative activity. But is it plausible that normative judgments consist (across the board) in decisions for others?

There is, I believe, an excellent preliminary case for this. As noted, we can express decisions for others in the context of delegated decision-making (and co-deliberation) using "should" statements ("you should get the burrito"). But all all-things-considered "should" statements have the same meaning, and so express the same state, regardless of the context in which they are uttered or accepted. That suggests that even in these "non-cooperative" contexts, the normative judgments these statements express are best understood as decisions. The additional intentions present in delegated decision-making and co-deliberation—e.g., the intention to communicate the conclusion of the relevant reasoning—are not constitutive of the *decision*. The decision is simply the *practical conclusion* which can be expressed in normative terms, and this presumably the same state we are in when making a private judgment about what someone should do, or overt judgments in contexts where the judgment stands no chance of being acted on.

---

[20]Note that if you concur with my judgment by making a decision for yourself, the state you agreed with must also have been a decision for you.

The proposal, however, raises a worry. The view assumes that we can coherently decide (in our technical sense) on something over which we know we have no control. Suppose I make a normative judgment about someone whose conduct I cannot affect. I judge that my husband should stop cheating on his partners, knowing he will not. Simple Plan Expressivism says that the judgment is a decision for my husband whose constitutive "point" is to affect his intentions. Since my normative *judgment* may be perfectly coherent (and apt), the plan expressivist must suppose that such knowably inefficacious *decisions* can be coherent (and apt). But is this plausible?

The first thing to note is that every plan expressivist must grapple with this problem of explaining the coherence of such decisions. As Gibbard observes, *first-personal* normative judgments—which the expressivist identifies with de se plans—are also often (knowably) inefficacious, though they may be exactly what the situation calls for. A binge alcoholic, knowing he will be driven by compulsion to drink on Saturday, can still judge—appropriately—that he should not. Gibbard claims that asked Wednesday for a decision whether to get drunk in his situation to come, he can decide not to, though "he can't make this hypothetical decision effective" (50). Consider also judgments about the past. My judgment that I should not have had that third drink last night can be appropriate even though the hypothetical plan that constitutes it has no chance of affecting my past conduct, as Gibbard acknowledges (50-51).

Since every plan expressivist must concede that some decisions (normative judgments) are inevitably inefficacious in this way, the fact that my view saddles us with many such plans is not a reason to favor Gibbard's Subtle View of third-personal judgments over Simple Plan Expressivism, which is the question I am concerned with in this article.[21] No doubt, the plan expressivist will emphasize that normative judgments are not intentions, but a *sui generis* state, which is subject to a distinct set of norms. For instance, while intention is subject to a Settling Condition (it is irrational to intend something you do not believe you can settle),[22] the norm on decision appears to be weaker: decisions, in our technical sense, should be such that nothing debars them from playing their designated output role *if circumstances arise where they are called upon to do so*. My husband *could* act on my judgment/decision (it is in his power to stop cheating on his partners), so it satisfies this weaker condition.[23]

Still, one might be skeptical: how could a state be apt if has no prospect—in a particular case—of fulfilling its normal output role? In response to this general worry, the plan expressivist, I believe, should lean on the observation that it is true of many types of states that they can be fully apt or warranted in their own terms—however these notions are understood—even when they cannot possibly play their official role.

Consider belief. In one sense, the aim of belief is to represent things as they are, so any true belief serves the purpose. But in another sense, the aim of belief is to guide action by channeling intentions from ends to means in ways that benefit the agent. That is arguably the purpose for which beliefs as a class exist, so a belief that represents the world accurately but fails to affect how the agent behaves has not played this type-defining role. But while the first sort of role determines a correctness condition for every token state, the second does not, though it may yield norms of other kinds. Beliefs are not inapt just because they stand no chance of guiding action.

[21]One might, of course, think that the plan expressivism's commitment to the aptness of such decisions makes plan expressivism untenable, but my aim in this paper is to find the most promising version of plan expressivism, not to defend the expressivist against the cognitivists.

[22]Bratman (1997).

[23]This output norm corresponds to an "ought implies can" constraint.

The second is closer to the conception of the role of a state I have been invoking: roughly, its "point" in our psychological economy, or what it does for agents in the course of deliberate action. The same, I say, goes for decisions. An inefficacious decision that is otherwise justifiable is like an inefficacious true belief: though it fails to do what states of that kind exist in order to do, it is not defective in itself.

Now let me show why Simple Plan Expressivism is not subject to the Second Objection, concerning motivational role. Simple Plan Expressivism, I submit, does not imply that anyone who judges that X should Φ is thereby motivated to promote X's Φing. To see why, consider again the functional role of third-personal judgments/decisions. According to the view sketched above, their role is to influence the intentions of someone else. Following Roth (2004), we can distinguish between the psychological subject of a decision—the individual who issues the decision—from the "executive subject" of the decision, the individual whose "job" it is to execute the decision. The executive subject of your first-personal decisions is you. But the executive subject your decision that I get the burrito is me. A decision generates a motivation only when it leads to an intention; but even when all goes well, a decision only leads to an intention in its executive subject. So while your decision for me that I get the burrito is poised to generate an intention and a motivation in me (and likely will if it is expressed), you might not care at all whether I fulfill your decision, nor are you rationally committed to caring, and so need not be motivated to help me get the burrito. After all the judgment/decision is meant to settle a practical question *for me*, not for you. Its functional role is *social*, extending "outside the head." Since the role of third-personal decisions is to motivate the executive subject of the decision rather than the psychological subject, Simple Plan Expressivism is consistent with the phenomenological datum that my judgment that Mary should leave her marriage does not imply an inclination on my part to help Mary leave.

This is a clear advantage of identifying the third-personal judgment that X should Φ with the decision *that* X Φ as opposed to the decision to do whatever one can to bring it about that X Φs. This nearby view would also avoid the objections to the Subtle View I set out in the first half of the paper, but it would be manifestly inconsistent with the motivational role of these judgments, as Ridge (2018) observes. It also shows why it is critical to identify third personal judgments with decisions that others do things rather than preferences or desires that others do things, since the role of desire and preference is always to motivate the *bearer* of the preference or desire to bring about the wanted state of affairs.

Simple Plan Expressivism, I believe, resonates with a compelling view of the role of normative judgment, according to which the "point" of these judgments—their raison d'être— is to control intentions, both intrasubjectively and intersubjectively, in light of reasons. As agents, we not only have the ability to step back from our own impulses and ratify or not the course of action they seem to favor, but also to guide the intentions of other agents in our milieux by deliberating from their point of view and ratifying or not courses of action for them. In the intrasubjective case, we often do a better job of acting on our reasons if we step back and reason from them, bringing the full range of our background knowledge to bear. Similarly, since it will often happen that others know more than we do and are less subject to some of the pathologies of decision than we are when we deliberate in our own case, it can be valuable for others to deliberate on our behalf, to express those decisions, and for us at least sometimes to act on them. The hypothesis I am advancing is that normative discourse exists so as to provide the resources for exercising this valuable capacity for intra- and intersubjective regulation of intention by decision for ourselves and others.

Contrast this with Gibbard's view, according to which the usefulness of third-personal judgments is to "rehearse for life" by settling in advance what we will do in remote counterfactual scenarios (Gibbard 2003: 52). Gibbard's view seems to misdescribe the well-functioning of third-personal judgments. It characterizes these judgments as most "at home" in first-personal, individual planning activity, rather than in the *social practices* in which they seem to do most of their distinctive work.

To sum up the argument of this section: Simple Plan Expressivism avoids the objections to Gibbard's conditional attitude view discussed in the first half of the paper (the problem of disagreement, the wrong subject matter objection, and the conflation objection), provides a uniform account of first and third-personal judgments, and can be supported by a compelling account of the point of normative judgment which is worth exploring in its own right. The two most pressing objections to the Simple View—the First and Second Objections—can be answered, as I have shown. This is an excellent case for preferring Simple Plan Expressivism to Gibbard's Subtle View, and for taking it seriously as a better articulation of the plan expressivist's main idea.

6 | RATIONAL CONSTRAINTS ON NORMATIVE JUDGMENT

In this section, I will show that whether we adopt Simple Plan Expressivism or Gibbard's Subtle View of third-personal judgments has implications beyond esoteric debates in metaethics. In particular, Simple Plan Expressivism entails rational constraints on normative judgment that do not follow from the Subtle View. This makes the choice between the two views consequential for normative ethics. I will suggest that these rational constraints are independently plausible. Hence, they ground an additional argument for Simple Plan Expressivism over the Subtle View.

Recall the burrito example: If you decide on my behalf that I'm to get the last burrito, you cannot coherently decide to get it yourself. Simple Plan Expressivism explains this by identifying judgments with decisions and noting that decisions are subject to a co-realizability constraint: it is irrational to decide that A Φ and B Ψ if it is impossible that [A Φ and B Ψ]. Given this principle, Simple Plan Expressivism entails a range of non-trivial coherence constraints on normative judgment. In particular, Simple Plan Expressivism entails that if you judge that X should Φ, you are committed to not interfering with X's Φing. The judgment that X should Φ is a decision that X Φ; the decision to interfere with X's Φing is a decision to do something incompatible with X's Φing. Since non-co-realizable decisions are incoherent, the corresponding normative judgments are likewise incoherent. So making the first commits you to not making the second.[24]

This is the element of the view that entails that rational egoism as standardly formulated cannot be coherently accepted. Rational egoism, on the standard formulation, says one should perform some action if and only if and because it maximizes one's self-interest. Suppose you and I are competing for some object. It would maximize my self-interest to acquire it for myself, and it would maximize your self-interest to acquire it for yourself. If I am an egoist, I am committed to the following judgments:

---

[24]Simple Plan Expressivism does not entail by this route that you are committed to promoting X's Φing if you judge that X should Φ, since the decision not to promote X's Φing is in general co-realizable with the decision that X Φ. The view thus vindicates Wallace's (2009) idea that the reasons of others provide us with reasons not to interfere with their rational choices without at the same to providing us with reasons to promote their ends.

(J1) the judgment that I should take the object for myself (because that would be best for me); and

(J2) the judgment that you should take the object for yourself (because that would be best for you)

Simple Plan Expressivism tells us that I must therefore have made the following decisions:

(D1) a decision for me that I take the object for myself

(D2) a decision for you that you take the object for yourself

But these decisions are incoherent in a straightforward sense. Just as a set of beliefs is coherent (consistent) only if they can all be true together, so a set of decisions is coherent only if they can all be realized together.[25] (D1) and (D2) are incoherent in this sense, so Simple Plan Expressivism entails that (J1) and (J2) are incoherent. And of course egoism is not the only problem case. The view entails the incoherence of any normative view which prescribes incompatible courses of action for individuals. More specifically, Simple Plan Expressivism entails that that the following *interpersonal joint satisfiability* principle must be true: if A should Φ and B should Ψ; then it is possible that [A Φ and B Ψ].[26]

This amounts to a formal refutation of egoism taking Simple Plan Expressivism as a premise. A possible reaction is to see this as a favorable, and possibly deep, implication of Simple Plan Expressivism. While egoism is now widely regarded as a coherent if mistaken view, given the history of the subject, we should see this consensus for what it is: a counsel of despair, an idea we've gotten used to only because we have not been able to do better. Philosophers throughout history have spilled much ink attempting to refute the egoist by disclosing some deep formal difficulty (Sidgwick 2019/1874, Moore 1903, Kant 1785, Medlin 1957, Baier 1958, Korsgaard 1996). We can all agree, as a substantive matter, that we should be altruistic to some extent. But *why should* one act, at least to some extent, for the sake of others and not only for one's own sake? The strongest possible answer would be a persuasive argument that would show egoist to be untenable or incoherent or otherwise unacceptable *by the egoist's own lights*: the sort of argument that would lead any reflective, truly rational egoist to revise his view in favor of some form of altruism.[27] Simple Plan Expressivism appears to give us such an argument, and if it did, that would be a good thing.

However, this line of reasoning is blunted by the fact that there is a view in the vicinity of rational egoism that is coherent on Simple Plan Expressivism: the view that everyone should *try* to do what is best for himself. It is not incoherent to decide that A try to Φ and that B try to Ψ even if it is impossible that A Φ and B Ψ, since it is always possible for each agent to try. The

---

[25]This argument resembles Medlin's (1957) argument against egoism which also takes non-cognitivism as a premise. Medlin argued that the egoist has incoherent preferences in situations like the one described. See Dreier (1996) for discussion.

[26]Most philosophers already accept the following *intrapersonal* joint satisfiability principle: if A should Φ and A should ψ; then it is possible that [A Φ and A ψ].

[27]Arguments in this vein include those of G.E. Moore (1903), Kant, and Medlin (1957).

egoist can simply retreat to this modified view. So there is no quick argument from Simple Plan Expressivism to a positive requirement of altruism.

But even if his is a disappointment, it is also a good thing. The coherence of these "trying" judgments provides a means of answering an objection to Simple Plan Expressivism and the rational constraints on normative judgment that it implies. After all, one might see the rational constraints entailed by the Simple Plan Expressivism not as a favorable implication of it, but as a serious blight. The idea that normative judgments must be jointly satisfiable sounds right when the judgments are expressed and the context is cooperative, as in delegated decision-making and co-deliberation. But it seems to clash with commonsense in other contexts. A burglar judges that the homeowner should call the cops. Cutting the phone lines, he knows, would preclude her from doing this. Is the burglar then incoherent if he proclaims that he should cut the phone lines? Intuitively not. So while Simple Plan Expressivism neatly explains the requirement of non-interference in delegated decision making, it appears to clash with intuition in other contexts.

But here we can invoke the "trying defense."[28] The trying defense says that what we often mean—or should mean—in contexts in which we say that people should do incompatible things is that each person should *try*.[29] When we say that people should be happy, or cure cancer, or write well, what we ought to say is that they should do what they can to be happy/cure cancer/ write well. Analogously, what the burglar ought to say is that the homeowners should try to call the police and that he should try to thwart them, since what is in their power is to try. If judgments like (J1) and (J2) strike us as coherent, this is because we are interpreting them as "trying" judgments.

Indeed, these observations may be turned into an argument for Simple Plan Expressivism. Note that on reflection, even in the absence of theory, there is something odd about (J1) and (J2) when they are not interpreted as "trying" claims. It would be fine to say that I should *try* to take the object for myself and you should try to thwart me. But it sounds odd to say outright that we each *should* take the object for ourselves since that's impossible. Suppose Alice and Bob come to you for advice and ask what they should do if they find themselves in such a situation. It would be odd for you in your capacity as advisor to say: "Well, everyone should do what's best for himself, so you, Alice, should take the object for yourself, and you, Bob, should take it for yourself." That sort of advice sounds not just false but incoherent. This suggests that the interpersonal joint satisfiability may well be true. If it is, it is a point in favor of Simple Plan Expressivism that it entails it.[30]

To summarize: The choice between Simple Plan Expressivism and the Gibbard's Subtle View consequential. If Simple Plan Expressivism is true, so is the interpersonal joint satisfiability principle. At first glance, this may seem to be problematic for Simple Plan Expressivism, because judgments like (J1) and (J2) initially appear to be coherent. However, on reflection, there

---

[28]The trying defense is discussed by Kalin (1970) and Dreier (1996) in response to Medlin's (1957) non-cognitivist argument for the incoherence of egoism.

[29]See Dreier (1996: 415) for a statement of this idea.

[30]One might object that the trying defense cannot succeed because sometimes trying is counterproductive. Consider the "paradox of hedonism": we often better succeed at achieving happiness if we do not deliberately seek it. The notion of "trying" can be understood in multiple ways. On one interpretation, to try to Φ is to aim at Φing. But "trying" to Φ can also mean doing what one can to Φ. If trying is interpreted in the latter way, the trying defense can survive the paradox of hedonism worry. To say that everyone should try to maximize their happiness, on this interpretation, is to say that everyone should do what they can to maximize their happiness. This needn't involve aiming at happiness. Thanks to an anonymous reviewer for raising this concern.

is something odd about (J1) and (J2). Simple Plan Expressivism explains this: (J1) and (J2) are incoherent because they consist in decisions for others which are subject to a co-realizability constraint. This is an additional argument for the Simple Plan Expressivism over both the Subtle View.

7 | CONCLUSION

I have argued that the best expressivist treatment of third-personal judgments of the form "X should Φ" identifies them with a kind of plan or decision *that X Φ*—a decision whose focus is someone else's conduct. This view of normative third-personal judgments—aside from being simple and elegant—avoids the objections to the Subtle View of third-personal judgments set out in the first half of the paper, resonates with a new and interesting functionalist moral psychology, and makes better sense of certain independently plausible constraints on normative judgment.

ORCID

Alisabeth Ayars https://orcid.org/0000-0002-5396-6269

REFERENCES

Baier, K. (1958). *The moral point of view: A rational basis of ethics.*

Beddor, B. (2018). Subjective disagreement. *Noûs*, forthcoming.

Bratman, M. E. (2013). *Shared agency: A planning theory of acting together.* Oxford University Press.

Bratman, M. E. (2006). Thinking how to live and the restriction problem. *Philosophy and Phenomenological Research*, *72*(3), 707–713.

Cullity, G. (2007). Thinking how to live ? Allan Gibbard. *The Philosophical Quarterly*, *57*(227), 308–311.

Dreier, J. (2015). Another world.

Dreier, J. (2009). Relativism (and expressivism) and the problem of disagreement. *Philosophical perspectives*, *23*, 79–110.

Dreier, J. (1996). Accepting agent centred norms: A problem for non-cognitivists and a suggestion for solving it. *Australasian Journal of Philosophy*, *74*(3), 409–422.

Egan, A. (2014). There's something funny about comedy: A case study in faultless disagreement. *Erkenntnis*, *79*(1), 73–100.

Egan, A. (2010). Disputing about taste.

Finlay, S. (2014). *Confusion of tongues: A theory of normative language.* Oxford University Press.

Finlay, S. (2010). What ought probably means, and why you canâ€™t detach it. *Synthese*, *177*(1), 67–89.

Finlay, S. (2009). Oughts and ends. *Philosophical Studies*, *143*(3), 315–340.

Gibbard, A. (2002). Normative and recognitional concepts. *Philosophy and Phenomenological Research*, *64*(1), 151–167.

Gibbard, A. (1983). A noncognitivistic analysis of rationality in action. *Social Theory and Practice*, *9*(2/3), 199–221.

Gilbert, M. (1996). *Living together: Rationality, sociality, and obligation.* Rowman & Littlefield Publishers.

Gregory, A. (2017). Are all normative judgments desire-like? *Journal of Ethics and Social Philosophy*, *12*(1), 29–55.

Kalin, J. In defense of egoism. Morality and Rational Self-Interest 1970:64–87.

Kant, I. Groundwork of the Metaphysics of Morals (1785 1785).

Korsgaard, C. (1996). *The sources of normativity.* Cambridge University Press.

Lewis, D. (1979). Attitudes de dicto and de se. *The philosophical review*, *88*(4), 513–543.

Manne, K. (2014). Internalism about reasons: Sad but true? *Philosophical Studies*, *167*(1), 89–117.

Medlin, B. (1957). Ultimate principles and ethical egoism. *Australasian Journal of Philosophy*, *35*(2), 111–118.

Richard, M. (2011). Relativistic content and disagreement. *Philosophical Studies*, *156*(3), 421–431.

Ridge, M. (2018). Expressivism and collectives. *Mind*, *127*(507), 833–861.

Ridge, M. (2013). Disagreement. *Philosophy and Phenomenological Research*, *86*(1), 41–63.

Ross, J., & Schroeder, M. (2013). Reversibility or disagreement. *Mind*, *122*(485), 43–84.

Roth, A. (2018). Interpersonal Obligation in Joint Action.

Roth, A. S. (2016). Intention, expectation, and promissory obligation. *Ethics*, *127*(1), 88–115.

Roth, A. S. (2014). Prediction, authority, and entitlement in shared activity. *Noûs*, *48*(4), 626–652.

Roth, A. S. (2004). Shared agency and contralateral commitments. *The Philosophical Review*, *113*(3), 359–410.

Schroeder, M. (2008). *Being for: Evaluating the semantic program of expressivism*. Oxford University Press.

Searle, J. R. Collective intentions and actions. Intentions in communication 1990;401:401.

Sidgwick, H. (2019). *The methods of ethics*. Good Press.

Silk, A. (2015). How to be an ethical expressivist. *Philosophy and Phenomenological Research*, *91*(1), 47–81.

Soames, S. Presupposition. Handbook of philosophical logic: Springer; 1989. p. 553–616.

Unwin, N. (2001). Norms and negation: A problem for Gibbard's logic. *The Philosophical Quarterly*, *51*(202), 60–75.

Wedgwood, R. (2006). The meaning of 'ought'. *Oxford studies in metaethics*, *1*, 127–160.

Worsnip, A. (2019). 'Ought'-contextualism beyond the parochial. *Philosophical Studies*, *176*(11), 3099–3119.