What Are We to Do? Making Sense of 'Joint Ought' Talk

Rowan Mellor and Margaret Shea (equal contributors)

I. Introduction

Sometimes it seems that two agents ought each to perform some act, but false that they ought both to perform it. Consider:

> **LONELY GRANDMOTHER** Amy and Brad's grandmother lives in a nursing home which rarely allows visits. Next Saturday is a rare occasion when visitors are allowed. Amy knows that Brad has booked a day at the spa, so will not visit. Brad knows that Amy has booked one, too, so will not visit. Due to her dementia, their grandmother becomes upset and confused if more than one relative visits on the same day.[1]

Amy ought to visit. For she knows that Brad is not going to visit, and that if she does not visit, nobody will. Amy's obligation is conditional – she ought to visit if and *only if* Brad is not going to visit – but crucially, the antecedent of the conditional is satisfied. As a matter of fact, Brad *is not* going to visit. Similarly – because he knows that Amy is not going to visit, and thus that, if he does not visit, nobody will – Brad ought to visit. So Amy and Brad each ought to visit. Nonetheless, 'Amy and Brad ought to visit' sounds false. If they both visit, their grandmother will fare worse than if she were left alone. So:

(1) Amy ought to visit

and:

(2) Brad ought to visit

but not:

(3) Amy and Brad ought to visit.

Similarly, sometimes it seems that two agents ought to perform some pair of acts, but false that each alone should perform his or her respective act:

[1] This is a variation on a case presented in (Shea forthcoming). (Jackson 1985) discusses a similar case.

> **SLICE & PATCH** Patient will die painlessly unless he receives a simple operation, which requires Dr. Slice to slice and Dr. Patch to patch. If Patient is sliced but not patched, or vice-versa, he'll die an agonizing death. Each surgeon knows that the other is going to play golf instead of doing their part.[2]

Slice ought not to slice because (as she knows) Patch is not going to patch; if she slices, Patient will die an agonizing death. By the same token, Patch ought not to patch. Thus the surgeons are conditionally obligated not to do their parts of the operation. Nevertheless, 'Slice and Patch ought to slice and patch' rings true. *Taken alone*, the best thing they can do is not do their parts. But *taken together*, the best thing for them to do is to do their parts. So:

(4) Slice and Patch ought to slice and patch.

but not

(5) Slice ought to slice.

and not

(6) Patch ought to patch.

These sets of 'ought' sentences – (1)-(3) and (4)-(6) – raise a question about the meaning of assertions of the form 'A and B ought to φ and ψ.' A first thought is that 'A and B ought to φ and ψ' serves to express the conjunction of claims 'A ought to φ' and 'B ought to ψ.' This is indeed one good interpretation of 'A and B ought to φ and ψ.' But it cannot be the only one. If it were, then fully competent, well-informed speakers would never take 'A and B ought to φ and ψ' to have a different truth-value from 'A ought to φ' and 'B ought to ψ.' But competent, well-informed speakers *do* sometimes take these to differ in truth-value. This is what LONELY GRANDMOTHER and SLICE & PATCH demonstrate: (1) and (2) are both true, while (3) is false; and (4) is true, while

[2] This case is due to (Estlund 2020, 211) and is discussed in further detail in (Mellor forthcoming). For cases with a similar structure, see *e.g.* (Gibbard 1965; Regan 1980).

(5) and (6) are both false. Our cases thus raise the question what *else* sentences of the form 'A and B ought to φ and ψ' can convey, besides 'A ought to φ' and 'B ought to ψ.'

This paper proposes that sentences of the form 'A and B ought to φ and ψ' can felicitously be used and interpreted by speakers to say something about what A and B *together* ought to do. On this alternative interpretation – what we call the *joint-ought interpretation* – 'A and B ought to φ and ψ' conveys not that A ought to φ and that B ought to ψ, but that *the plurality* A and B ought to φ and ψ.[3] For ease of exposition, when we wish to discuss a sentence's joint-ought interpretation, we will refer to it as a *joint-ought claim*.[4] Importantly, the joint-ought claim 'A and B ought to φ and ψ' can differ in truth-value from 'A ought to φ' and 'B ought to ψ.'[5]

The paper proceeds as follows. In Section II, we argue that what several agents *together* ought to do can come apart from what they *individually* ought to do.[6] [7] This explains why the joint-ought claim 'A and B ought to φ and ψ' can have a different truth-value from 'A ought to φ' and

[3] A third available interpretation of 'A and B ought to φ and ψ' must be mentioned. The sentence can be used to say that A and B each ought to perform the pair of acts φ and ψ. This interpretation is not our concern. The details of SLICE & PATCH are such that the case does not give rise to the intuition that either doctor should both slice and patch, while LONELY GRANDMOTHER features only one act-type. So neither sentence (3) nor (4) is plausibly interpreted in this way. We set this third interpretation aside in what follows.

[4] (Pinkert 2014) introduces a similar term, "joint obligation."

[5] On some semantic models for 'ought' the joint-ought interpretation of 'A and B ought to φ and ψ' is truth-conditionally equivalent to a conjunction, while on others it is not. We discuss this further in §III-IV.

[6] (Jackson 1987, 1988), (Parfit 1988), (Smith 2009), (Pinkert 2014), (Dietz 2016), and (White forthcoming) all defend versions of this claim. Others, such as (Schwenkenbecher 2019, 2021), defend the idea that some obligations attach to *pluralities* of agents, but deny that these obligations can diverge from the obligations of the individuals in the plurality.

Our proposal also evokes discussions of moral dilemmas (Williams 1973) and Professor Procrastinate cases (Jackson and Pargetter 1986). These are cases in which it appears that *a single person* ought to perform each of two acts but not the pair of them; and in which it appears that *a single person* ought to perform a pair of acts, but not each of them. Our cases, by contrast, are ones in which it appears that each, but not both, of two acts involving *two different persons* should be performed; and in which it appears that both, but not each, of two acts involving *two different persons*, should be performed. Note that the significance of this distinction will depend in part on one's conception of diachronic agency. On a Humean view, according to which an agent *just is* a time-slice that has control over all its internal elements, Professor Procrastinate cases are multi-agent cases, since they concern what different time-slices of a person (thus, what different *agents*) should do.

[7] (Mellor forthcoming) argues that this provides the best explanation of the intuitions prompted by SLICE & PATCH.

'B ought to ψ.' Sections III and IV explain how joint-ought claims can be modeled semantically, in a way which allows for the differences in truth-value we emphasize. We demonstrate how this can be accomplished within the framework of Mark Schroeder's relational semantics for 'ought,' and within the framework of Angelika Kratzer's modal semantics for 'ought.' Finally, Section V addresses an objection to our proposed interpretation of sentences like (3) and (4), which draws upon the widely-acknowledged distinction between 'deliberative' and 'evaluative' ought sentences.

II. Oughts, Options, and Outcomes

An individual-ought claim, of the form 'A ought to φ,' has an essentially action-guiding character;[8] it marks out a course of action as 'to be done' by the agent A.[9] The joint-ought claim 'A and B ought to φ and ψ,' we propose, has the same action-guiding character: it marks out a *series of acts*, φ-and-ψ, as to be done by the *plurality* of agents, A and B.[10] The joint-ought claim 'A and B ought to φ and ψ' is *not* the claim that φ-and-ψ is to be done by a *group agent* comprising A and B.[11] That can be construed as an individual-ought claim, 'G ought to φ and ψ,' where 'G' denotes the group agent comprising the natural persons A and B. The subject-position of a joint-ought claim is occupied not by any one agent (group or otherwise), but by *several agents*.[12]

[8] Thus, as we use the term, "individual-ought claims" are claims which have an individual agent as their subject and which feature 'ought' in its *deliberative* sense (*see* §V).

[9] (Gibbard 2003). Many philosophers cash 'to-be-doneness' out in terms of conclusive reason: the act which is to be done by an agent is that act which she has conclusive reason to perform.

[10] If to-be-doneness is cashed out in terms of conclusive reason, then our view implies that there are *joint reasons* which bear irreducibly on pluralities of agents.

[11] For a view of this sort, see e.g. (Killoren and Williams 2013). See (Mellor forthcoming) for an extended argument that joint-ought claims do a better job of accounting for our intuitions about cases like SLICE & PATCH than ought-claims attaching to group agents.

[12] (Pinkert 2014, p. 189) makes this same distinction.

Furthermore, the acts implicated in a joint-ought claim need not constitute a 'joint activity' in the sense popularized by authors such as Michael Bratman (2014) and Margaret Gilbert (2013). That is, we leave it open that A and B could succeed in doing what they jointly ought to do merely by each performing an individual act, without forming a 'shared intention' or 'joint commitment' (or whatever else might be required for joint agency).

In this section, we show that, granting a plausible and widespread understanding of ought-claims and their relation to options, what several agents *jointly* ought to do can diverge from what they *individually* ought to do. In turn, the joint-ought claim 'A and B ought to φ and ψ' can differ in truth-value from the individual-ought claims 'A ought to φ' and 'B ought to ψ.'

Whether a course of action is 'to be done' by an individual agent depends on whether it is one of her options, and on how it compares to her other options.[13] How exactly to identify and individuate an agent's options is a fraught question. Broadly speaking, however, an agent's options track what is *under her control*. Plausibly, an agent A's option-set is the set of actions which it is within her power to perform,[14] where this power can be roughly characterized in counterfactual terms: it is within A's power to φ if (i) A would φ if she intended to, and (ii) A can intend to φ.[15] Similarly, whether a course of action is to be done by a plurality of agents depends on whether it is one of their *joint options*, and on how it compares to their other joint options. We propose that,

[13] See e.g., (Portmore 2019, chap. 1)

[14] (Portmore 2019, chap. 2).

[15] This counterfactual characterization of an option-set is intended as a heuristic, not a reductive definition. We assume that an agent's options overlap significantly, *but not perfectly*, with what she would do if she intended (where she can so intend). It is a familiar point that giving a reductive analysis of the notion of an option, or ability, in these counterfactual terms raises problems. For one, the characterization of an option is circular, since clause (ii) invokes what an agent *can* intend. Removing clause (ii) would eliminate this circularity, but over-generate options: an agent who would ϕ if she intended to, but who is incapable of forming the intention to ϕ, would be counted as having the option ϕ. For another, the characterization under-generates abilities in cases like that of J. L. Austin's (1956) self-reproaching golfer, who intends to make the putt, and misses, yet knows that he *could* have made it.

on analogy to individual options, a plurality's joint options are those act-series which it is within their joint power to perform. Again, this can be roughly characterized in counterfactual terms: an act-series χ is within A and B's joint power if (i) A and B would χ if they were to intend the components of χ, and (ii) A and B can intend the components[16] of χ.[17] For example, it is within Maggie and Rowan's joint power to tango if they would tango if Maggie were to intend to dance one part (and can intend this) and Rowan were to intend to dance the other (and can intend this).

The counterfactual characterization of individual and joint option-sets illustrates that what A and B *together* can do may differ from what A can do plus what B can do. That is: a plurality's joint options can diverge from the combinations of their individual options. For whether some agent(s) have the option to perform some acts depends on what happens in the nearest possible world in which they intend to perform these acts, and the nearest possible world in which A intends to φ and B intends to ψ may be neither the nearest possible world in which A intends to φ nor the nearest possible world in which B intends to ψ. Given the connection between options and 'ought's, this marks a first way that joint-ought claims and pairs of individual-ought claims can come apart: a joint option can fail to decompose into individual options, and individual options can fail to compose into a joint options.

Some examples illustrate the point. Suppose that Maggie and Rowan are anarchic rowing partners. If either were to attempt to move the boat forward, the other would immediately stop

[16] Just what it takes for a plurality of agents to intend the components of some act-series can vary depending on how *structured* this act-series is. Note that some joint-ought claims say not only that a plurality of agents ought to perform some act-series, but, more specifically, that certain members of this plurality ought to perform certain acts (*see* §III). In these cases, intending the components of the act-series will require certain members of the plurality to intend certain of the acts in the series.

[17] Again, this is intended as a rough heuristic, not a reductive analysis. If it *were* intended as a reductive analysis, it might face additional problems (beyond those discussed in n. 15) when applied to pluralities. (Pinkert 2014), for instance, argues that in order for χ to be a joint option for A and B, there must be a possible combination of the components of χ *which is salient to them*. For further discussion of the notion of joint ability, see e.g. (Estlund 2020, chap. 12; Pinkert 2014; Schwenkenbecher 2021, chap. 3; Collins 2019, chap. 3; Spiekermann 2021).

rowing, and the boat would spin in a circle. Neither Maggie nor Rowan possesses the option of rowing the boat forward: forward motion requires rowing from both of them, and the nearest possible world in which either intends to row the boat forward is a world where the other does not. However, Maggie and Rowan do possess the *joint* option of rowing the boat forward: in the nearest possible world in which they *both* intend to row it forward, they succeed. Here, Maggie and Rowan's joint option does not decompose into individual options possessed by either of them. So while it could be true that they *jointly* ought to row the boat forward, it could not be true that either individually ought to do so.

Similarly, there are cases where several agents' individual options fail to compose into a joint option. Suppose that Maggie and Rowan are locked in a cell together when their jailor makes them an offer. She is willing to let one of them go free. But there is a catch: if they both try to leave, then she will rescind her offer and keep both captive. As it happens, Maggie and Rowan both like cell-life, so neither plans on leaving.[18] Here, Maggie and Rowan both have the individual option to leave. The nearest possible world in which Maggie intends to leave is a world in which Rowan *does not*, and thus a world in which she succeeds; and the same goes for Rowan. However, Maggie and Rowan lack the joint option of both leaving, since in the nearest world where they both intend to leave, they fail. So, while it might be that they each individually ought to leave, it cannot be that Maggie and Rowan jointly ought to leave.

There is also a second way that joint-ought claims diverge from pairs of individual-ought claims. What some agent(s) ought to do depends not only on what options they have, but also (at least in part) on the comparative value of their options. The value of an option can be thought of as the value of its (expected) outcome: an option is better if it has a better expected outcome, and

[18] This example is adapted from (Cohen 1983).

worse if its expected outcome is worse.[19] Like options, outcomes can be characterized counterfactually: the outcome O of action A is the world that would be actual were A performed.[20] As this characterization makes clear, the outcome of A's φ-ing *and* B's ψ-ing is distinct from the outcome of A's φ-ing, and the outcome of B's ψ-ing. For the nearest possible world in which A φ's *and* B ψ's may be neither the nearest world in which A φ's nor the nearest world in which B ψ's. So, even where a plurality's joint options do not diverge from the list of combinations of their individual options (the possibility discussed above), the comparative value of these options can differ, in light of their different outcomes. The outcome of A's φ-ing *and* B's ψ-ing might be much better or worse than the outcomes of A's φ-ing, and of B's ψ-ing.[21]

This is precisely what we find in LONELY GRANDMOTHER and SLICE & PATCH. In SLICE & PATCH, the surgeons each have two individual options: Slice can slice or not-slice, and Patch can patch or not-patch. There are four possible combinations of these individual options, and these exhaust the surgeons' joint options. That is: their joint options are identical to the conjunctions of their individual options. However, the comparative value of these joint options diverges from that of the individual options which they comprise. The best of the surgeons' joints options is clearly slicing and patching, since this has the best expected outcome: in the nearest world in which Slice slices and Patch patches, Patient survives. But slicing is Slice's worst option, since the nearest world where Slice slices is one where Patch does not patch; and patching is Patch's worst option, since the nearest world where Patch patches is one where Slice does not slice. So, while Slice and

[19] This does not assume consequentialism. For one, we define 'outcome' very broadly. For another, one can hold that the expected outcome of an option plays *some* role in determining its value, and therefore in determining what an agent ought to do, without holding that it is the *sole* determinant.

[20](Portmore 2011, 120).

[21] (Jackson 1988) and (Smith 2009) both make versions of this observation. *See*, in particular, Smith's discussion of "non-distributive sub-optimality" (37-39).

Patch jointly ought to slice and patch, it is false that Slice ought to slice and that Patch ought to patch.

Similarly, in LONELY GRANDMOTHER, Amy and Brad each have two individual options: Amy can visit or not visit, and Brad can visit or not visit. The four possible combinations of these individual options exhaust their joint options. But the comparative value of these joint options differs from that of the individual options which they comprise. Amy's best option is visiting, since the nearest possible world in which she visits is one in which Brad does not visit, and her grandmother benefits. By the same reasoning, Brad's best option is visiting. But *both* visiting is Amy and Brad's worst joint option, since it has the worst expected outcome: in the nearest world in which Amy visits and Brad visits, their grandmother is upset and confused. So, while each individually ought to visit, it is false that Amy and Brad jointly ought to visit.

III. Relational Semantics for Joint-Ought Claims

We have argued that what several individuals *jointly* ought to do can differ from what they *individually* ought to do. Accordingly, the joint-ought claim 'A and B ought to φ and ψ' can differ in truth-value from the pair of individual-ought claims 'A ought to φ' and 'B ought to ψ'. This raises the question whether standard semantic analyses of 'ought' can model this difference in truth-values. We consider two opposing analyses – one *relational*, one *modal* – and show that each of them can do so.

There is a clear distinction, in ordinary language, between joint and individual satisfaction of a predicate.[22] Predicates which attach to a plurality, and fail to distribute across a conjunction, are pervasive: 'Kim and Paul outweigh Mark' can be true, while 'Kim outweighs Mark' and 'Paul

[22] (McKay 2006, 8). While standard first-order logic is blind to this distinction, some philosophers have undertaken to develop plural logics and semantic accounts of *plural predication*, which can capture it. *See* (McKay 2006; Linnebo 2022; Oliver and Smiley 2016; Rayo 2007).

outweighs Mark' are false. So are predicates which attach to each of a set of individuals, but fail to agglomerate across a conjunction, and so do not attach to the plurality made up of those individuals: 'Kim is of one gender' and 'Paul is of one gender' can both be true, while 'Kim and Paul are of one gender' is false.[23]

When it is not analytic that a predicate P applies to a plurality of things if and only if it applies to each of them, P is *non-distributive*.[24] There is a striking similarity between the behavior of such predicates, and the behavior of 'ought' in joint-ought claims like (4) 'Slice and Patch ought to slice and patch,' which can be true even as (5) 'Slice ought to slice' and (6) 'Patch ought to patch' are false.[25] A first thought, then, is that 'ought' can be modeled as a non-distributive predicate.[26]

This proposal dovetails with a *relational approach* to the semantics of 'ought', which construes 'ought' as a two-place predicate relating an agent to something else.[27] Different views in this camp make different proposals about the second argument-place of this relation: on some, it is an act,[28]

---

[23] (McKay 2006, 8).

[24] (McKay 2006) distinguishes 'non-cumulative' predicates (like 'of one gender') from 'non-distributive' predicates (like 'outweigh'). For ease of exposition, we follow (Linnebo 2022) and roll the two notions into one biconditional.

[25] There is also a striking difference between what one might call 'strict' non-distributiveness and the distributive failure associated with 'ought' which we emphasize. Strict non-distributiveness is that phenomenon wherein one and the same predicate attaches to a plurality without necessarily attaching to its individual members, such that abRcd is semantically distinct from (aRcd and bRcd). For example, 'outweigh' in 'Kim and Paul outweigh Mark and Marcia' is strictly non-distributive: the sentence does not entail 'Kim weighs more than Mark and Marcia' or 'Paul weighs more than Mark and Marcia.' Supposing that 'ought' is a dyadic predicate, it, too, is strictly non-distributive. For example, 'Slice and Patch ought to slice and patch' does not entail 'Slice ought to slice and patch' or 'Patch ought to slice and patch.' But that is not the distributive failure associated with joint-ought sentences which occupies us. Rather, we are out to model the fact that it can be true that A and B ought to perform an act-series χ, without its being true that A or B ought to perform her part of χ. For example, 'Slice and Patch ought to slice and patch' does not entail 'Slice ought to slice' or 'Patch ought to patch.' Therefore, our subsequent references to a dyadic predicate R's "non-distributiveness" are references to the following failures of entailment: abRcd does not entail aRc or bCd, and aRc and bRd do not entail abRcd.

[26] (Pinkert 2014) develops a version of this proposal (*see* nn. 29 and 32). (Smith 2009) also discusses non-distributive predication.

[27] *See* n. 30.

[28] (Schroeder 2011).

on others, an agent-centered proposition.[29] We will focus on Mark Schroeder's (2011) view that 'ought' relates an agent and an action.[30] On this view, a sentence of the form 'A ought to φ' is modeled as 'A *ought* φ,' where the latter is true iff A stands in the *ought* relation to the act of φ-ing.

In expounding the idea that 'ought' is a two-place predicate, Schroeder assumes that both argument-places will be occupied by singly referring terms, picking out an agent and an act respectively.[31] But his account can easily accommodate joint-ought claims, by allowing both argument-places of the 'ought' relation to be occupied by *plurally referring terms*. This supplemented relational semantics has it that, just as in an individual-ought claim, 'ought' relates a single agent to a single act or act-series, so, in a joint-ought claim, 'ought' relates several agents to an act-series.[32] [33] This semantic model demystifies our key claim that the joint-ought claim 'A and B ought to φ and ψ' can differ in truth-value from the individual-ought claims 'A ought to φ'

[29] (Wedgwood 2006; Broome 2013). *See* (Pinkert 2014) for a Broomean modeling of obligations shared by a plurality (discussed further in n. 32).

[30] On Schroeder's (2011) view, 'ought' is systematically ambiguous between a *deliberative* sense and an *evaluative* sense. It is 'ought' in its deliberative sense which he claims functions as a dyadic relation. The distinction between deliberative and evaluative senses of 'ought' is discussed further in §IV.

[31] This is explained by the fact that he only discusses individual-ought claims of the form 'A ought to φ' – which choice likely reflects the prevalent assumption that ought-claims of the form 'A and B ought to φ and ψ [respectively]' are always equivalent to the conjunction of 'A ought to φ' and 'B ought to ψ.' (Shea forthcoming) argues that there is an important distinction between the two, which a plausible analysis of 'ought' must recognize; and points out that prominent meta-ethical theories, including Gibbard's (2003) plan-expressivism, fail to do so.

[32] (Pinkert 2014) makes a similar proposal about obligations attached to pluralities of agents. Adopting Broome's (2013) semantics for 'ought,' on which 'ought' is a two-place predicate which relates an agent to an agent-centered proposition, Pinkert argues that these obligations can be understood as two-place, non-distributive predicates, relating a plurality of agents to a multi-agent-centered proposition.

[33] On this semantic view, the sentence 'A and B ought to φ and ψ' is, on its joint interpretation, not truth-conditionally equivalent to the conjunction of the sentences 'A ought to φ' and 'B ought to ψ.' For the joint-ought claim 'A and B ought to φ and ψ' is modeled as 'AB *ought* φψ,' while the individual-ought claims 'A ought to φ' and 'B ought to ψ' are modeled as 'A *ought* φ ∧ B *ought* ψ.' As we argue in this section, these have different truth-conditions.

and ‘B ought to ψ.’ For, rendered relationally, the joint-ought claim does not entail, nor is it entailed by, the two individual-ought claims.

Consider first that, on our supplemented relational semantics, ‘A ought to φ’ and ‘B ought to ψ’ do not entail ‘A and B ought to φ and ψ’. Modeled relationally, this failure of entailment is pellucid: that A stands in the *ought* relation to φ and B stands in the *ought* relation to ψ, does not entail that A and B *together* stand in the *ought* relation to φ and ψ. There is no rule of inference:

aRb
cRd
∴
acRbd

The fact that Alex cleans the house faster than Beverly and the fact that Christie cleans the house faster than Dennis do not entail that Alex and Christie clean the house faster than Beverly and Dennis. For Beverly and Dennis may be better collaborators. On the proposed relational semantics, ‘ought’ behaves like ‘clean faster than’ in cases like LONELY GRANDMOTHER: Amy and Brad each individually stand in the *ought* relation to visiting, but Amy and Brad do not *together* stand in the *ought* relation to the combination of acts [visit, visit].[34] Thus, (1) ‘Amy ought to visit’ and (2) ‘Brad ought to visit’ are true while (3) ‘Amy and Brad ought to visit’ is false.

Second, on our supplemented relational semantics, the joint ‘ought’ sentence ‘A and B ought to φ and ψ’ entails neither ‘A ought to φ’ nor ‘B ought to ψ.’ Again, once modeled relationally, this failure of entailment is pellucid: that A and B together stand in the *ought* relation to φ and ψ entails neither that A stands in the *ought* relation to φ, nor that B stands in the *ought* relation to ψ. For there is no rule of inference:

acRbd

[34] The example happens to be one in which φ and ψ are the same act-type, with the notable result that the act-series implicated in joint-ought claim (3) does not require that Amy versus Brad perform any particular act in that series. *See* discussion of implicit ‘respectively’ clauses below.

∴
aRb v cRd

From the fact that Alex and Christie together clean the house slower than Beverly and Dennis, it does not follow that Alex cleans it slower than Beverly or that Christie cleans it slower than Dennis; Alex and Christie might be faster individual cleaners, but worse collaborators, than Beverly and Dennis. On the proposed relational semantics, 'ought' behaves like 'clean slower than' in cases like SLICE & PATCH, where Slice and Patch together stand in the 'ought' relation to slicing and patching, but Slice individually does not stand in the 'ought' relation to slicing, and Patch individually does not stand in the 'ought' relation to patching. Thus, (4) 'Slice and Patch ought to slice and patch' is true while (5) 'Slice ought to slice' and (6) 'Patch ought to patch' are false.

In short, there are many familiar dyadic predicates R which can take singular or plural arguments, and which are non-distributive in the following sense:[35] they do not license inferring aRb or cRd from acRbd, nor acRbd from aRb and cRd. On the supplemented relational semantics we propose, 'ought' is just another such predicate.[36]

An objection to this proposal suggests itself, however. Notice that some joint-ought sentences say something about *which* of the implicated plurality of agents ought to perform *what component* of the implicated plurality of acts. For example, (4) 'Slice and Patch ought to slice and patch' says *not* that Slice and Patch should perform the combination of acts [Patch slices, Slice patches], but that Slice and Patch should perform the combination of acts [Slice slices, Patch patches]. In other words, (4) seems to say that Slice and Patch ought to slice and patch

[35] *See* n. 25.

[36] To be clear, our proposal is that individual-ought sentences of the form 'A ought to φ' and joint 'ought' sentences of the form 'A and B ought to φ and ψ' feature one and the same relation, *ought*. We do *not* propose that the natural language word 'ought' is ambiguous between two distinct relations, one of which takes a singular argument (an agent) in its first argument-place and the other of which takes a plural argument (a plurality of agents) in its first argument-place.

respectively.[37] The problem is that, as a general rule, non-distributive predicates cannot be modified by a respectively clause.[38] For example, the sentence:

(7) Sue and Karen hired Bob and Tim[39]

admits of a reading on which 'hire' functions as a non-distributive predicate:

(7_1) Sue and Karen *together* hired Bob and Tim

And of a separate reading on which 'hire' is modified by a respectively clause:

(7_2) Sue and Karen hired Bob and Tim *respectively*

But there is no admissible reading of (7) on which 'hire' is *both* non-distributive *and* modified by a respectively clause:

(7_3) #Sue and Karen *together* hired Bob and Tim *respectively*

If this general rule held universally – such that there were *no* English sentences in which a non-distributive predicate seemed to be modified by an implicit 'respectively'[40] – that would constitute a serious challenge to the proposed relational analysis of joint-ought claims. For joint-ought claims like (4) 'Slice and Patch ought to slice and patch' seem to be modified by an implicit 'respectively,' and, granting the proposed analysis, 'ought' functions as a non-distributive predicate in these same claims.

In response to this worry, we observe that ordinary language provides examples of perfectly intelligible sentences which feature a predicate that is non-distributive *and* appears to be

[37] The semantics of 'respectively' constructions have received considerable attention from linguists. For discussion, *see e.g.* (Cho 2013; Chaves 2012; Gawron and Kehler 2004).

[38] Here is another worry in the same vein. We have proposed that, in joint-ought claims, *ought* relates plural arguments. But in joint-ought claims like 'Slice and Patch ought to slice and patch,' *ought* relates Slice to slicing and Patch to patching, *and not the other way around*. It is thus as though *ought* relates ordered arguments. But the notion of an ordered plurality is utterly mysterious.

[39] This example is taken from (Chaves 2012, 298).

[40] (Cho 2013, 282) claims that "'respectively' cannot co-occur with collective words."

modified by a 'respectively' clause – just as, on the relational view, (4) does. Consider, for instance:

(8) Abbott and Costello get big laughs for their straight face and clowning

Abbott plays the straight man and Costello the clown, and never *vice versa*: it is this assignment of roles which is comedically effective. So, on its true interpretation, this sentence says that Abbott and Costello stand in the 'get big laughs for' relation to the series of acts comprising *Abbott's* straight-facing and *Costello's* clowning (not Costello's clowning, and Abbott's straight face). However, thus interpreted, the sentence does not entail that Abbot, taken in his own right, gets big laughs for his straight face; nor that Costello, taken in his own right, gets big laughs for his clownishness. In the absence of the other, each component of the famous double-act might fall flat. Thus, 'get big laughs for' in (8) is non-distributive.

We take no stand on whether (8) is, in fact, best modeled as a sentence in which a non-distributive predicate is modified by an implicit 'respectively' clause. Linguists may disagree on how to model it. Our point is simply that this sentence exhibits the same combination of qualities which – granting the relational analysis of 'ought' – the joint-ought claim (4) 'Slice and Patch ought to slice and patch' exhibits. On this analysis, 'ought' in (4) is a non-distributive predicate, and nonetheless (4) expresses the claim that Slice and Patch stand in the 'ought' relation to the act-series comprising *Slice's* slicing and *Patch's* patching (as opposed to Patch's slicing and Slice's patching). The intelligibility of sentences like (8) suggests that this is not an entirely outré posit.

IV. Modal Semantics for Joint-Ought Claims

Of course, it may be that 'ought' should not be construed as a predicate at all, but instead as a necessity-like propositional operator quantifying over possible worlds. On this modal approach, 'A ought to φ' is modeled as OUGHT[A φ's], where this is true if and only if A φ's in the top-

ranked worlds over which 'ought' quantifies. On the simplest modal semantics, the operator OUGHT ranges over some fixed totality of possible worlds: 'A ought to φ' is true iff A φ's in the best worlds in this fixed domain. This simple semantics cannot draw the distinction in truth-value we emphasize, between the joint-ought claim 'A and B ought to φ and ψ' and the pair of individual-ought claims, 'A ought to φ' and 'B ought to ψ.' For there is no distinction between (i) its being the case that A φ's and B ψ's in the best of a fixed domain of possible worlds, and (ii) its being the case that A φ's in the best of this same domain of worlds and that B ψ's in the best of these same worlds.[41]

However, few proponents of the modal approach defend this simple semantics. On the most prominent modal semantics for 'ought' – Kratzer's (1981) *contextualist* view – OUGHT ranges over *different* domains, depending on the sentence in which 'ought' occurs and the context in which that sentence is asserted and evaluated.[42] In this section, we demonstrate that Kratzer's contextualist semantics for 'ought' can readily accommodate the difference we emphasize, between the truth-value of the joint-ought claim 'A and B ought to φ and ψ' and the pair of individual-ought claims 'A ought to φ' and 'B ought to ψ.'

Kratzerian contextualists hold that 'ought' is context-sensitive to two contextual parameters: a *modal base* and an *ordering source*.[43] The modal base comprises a set of propositions held fixed

[41] Take LONELY GRANDMOTHER. On the simplest modal semantics, (1) 'Amy ought to visit' is true if and only if Amy visits in the best of a fixed domain of worlds *D*, while (2) 'Brad ought to visit' is true if and only if Brad visits in the best of *D*. So, if (1) and (2) are both true, then the joint-ought interpretation of (3) 'Amy and Brad ought to visit' must also be true – since, on the simplest modal view, that joint-ought claim is true if and only if Amy and Brad visit in the best worlds in *D*. But this just what we deny: we hold that, though (1) and (2) are true, the joint-ought claim (3) 'Amy and Brad ought to visit' is false. *See* n. 44.

[42] (Kratzer 1981; 2012) develops a contextualist semantics for modals in general. *See* Chrisman (2016) and (Worsnip 2019) for applications of this semantics to 'ought' in particular.

[43] We focus on Kratzerian contextualism because it is widely regarded as the orthodox modal semantics. But other modal accounts, *e.g.* a 'contrastivist' semantics for 'ought' (Jackson 1985; Finlay and Snedegar 2014), may also have the resources to draw the distinction in truth-value we emphasize.

(treated as given) by speakers at a given context *c*, which restrict the domain of worlds over which OUGHT quantifies at that context (*i.e.*, the domain of so-called 'live' worlds). At *c*, OUGHT quantifies over only those worlds which are consistent with the modal base. The ordering source is a contextually-given set of standards which determines a ranking of the live worlds at *c*. OUGHT[*p*] is true at a context *c* if and only if, given the ranking of live worlds determined by *c*'s modal base and ordering source, the top-ranked worlds are worlds in which *p* holds.

We propose that there is a systematic difference between the modal bases relative to which joint-ought claims of the form 'A and B ought to φ and ψ' are evaluated, and those relative to which the individual-ought claims 'A ought to φ' and 'B ought to ψ' are evaluated. This explains their difference in truth value.[44] The context-shifts we identify are systematic in that they are all (we will argue) explained by the same pragmatic rule.

Without some pragmatic supplementation of the kind we go on to propose, Kratzer's semantics for 'ought' would conflict with an intuitive normative observation: *viz.*, the fact that an agent will not do something which she could do cannot in itself make it false that she ought to do that thing. For example, suppose that a billionaire stubbornly asserts that he will never file his tax returns. Intuitively, this cannot, in and of itself, make it false that he ought to file the returns. Agents cannot escape their obligations simply by foreseeably failing to comply with them.[45] Typically, however,

[44] Note that, on the semantic view discussed in this section, the joint-ought interpretation of the sentence 'A and B ought to φ and ψ' is truth-conditionally equivalent to the conjunction of the sentences 'A ought to φ' and 'B ought to ψ.' The former is modeled as OUGHT[A φ's and B ψ's], and the latter as conjunction of OUGHT[A φ's] and OUGHT[B ψ's]. These are truth-conditionally equivalent: relative to a single context, these always have the same truth-value. On the view discussed here, the difference in truth-values between the joint-ought interpretation of 'A and B ought to φ and ψ' and the conjunction of sentences 'A ought to φ' and 'B ought to ψ,' arises at the level of *pragmatics*. When speakers interpret sentences of the form 'A and B ought to φ and ψ' as joint-ought claims, this view proposes, they evaluate them relative to a different context from that relative to which they evaluate 'A ought to φ' and 'B ought to ψ.'

[45] As remarks like 'He should, but he won't' attest. (Shea forthcoming) emphasizes this point.

propositions which are general knowledge at a context are included in that context's modal base. It is general knowledge that our billionaire will not file his tax returns. So *if* speakers treated this assertion like a garden-variety proposition, holding it fixed in the conversational context at *c*, then 'Mr. Billionaire ought to file his tax returns' would be false at *c*. For there is no world consistent with the modal base at *c* in which Mr. Billionaire files his tax returns. But this is *not* what speakers do. Again: Mr. Billionaire's assertion that he will never file does not make it false that he ought to do so.

Kratzerians, we suggest, should take this as evidence that speakers use something like the following pragmatic rule:

> **RULE:** For any ought-claim of the form 'A ought to φ', and for any context *c* relative to which it is evaluated, if the modal base at *c* includes the proposition that A can φ, then it must exclude the proposition that A will not φ.

In other words, the assertion that the subject of an ought-claim can perform the act(s) implicated in it, tells speakers to keep worlds in which this subject fails to perform these act(s) *alive*: it serves as a constraint against narrowing the set of live worlds. The proposed behavior of 'A can φ' is unusual. For the addition of a proposition to a modal base generally *kills* worlds, having shrunk the domain of worlds over which OUGHT quantifies at that context. Yet the proposed behavior is not without precedent. Epistemic modals, for example, constrain the elimination of live worlds. 'Brad might quit' tells speakers to keep worlds in which Brad quits alive. Our proposal is that 'A can φ' does something similar, in a specific kind of context. At any context of evaluation for 'A ought to φ,' 'A can φ' tells speakers to keep worlds in which A φ's alive.

This pragmatic rule explains our judgments about LONELY GRANDMOTHER and SLICE AND PATCH. Ought-claims (1)-(3) and (4)-(6) are each evaluated relative to different contexts. For the

pragmatic rule guides speakers to look to the occupant(s) of the subject-position of an ought-claim in determining which worlds to treat as live, and each of ought-claims (1)-(3) and (4)-(6) has a different agent (or set of agents) in its subject-position.

Take LONELY GRANDMOTHER. It is clear from the description of the case that Amy and Brad each can visit, but will not. The case has four possible outcomes, and we can assume an ordering source which ranks them as follows:

TIED BEST	**w_1** Amy visits and Brad doesn't visit
TIED BEST	**w_2** Brad visits and Amy doesn't visit
2ND BEST	**w_3** Amy doesn't visit and Brad doesn't visit
WORST	**w_4** Amy visits and Brad visits

The pragmatic rule we have identified predicts that (1) 'Amy ought to visit' is evaluated relative to a modal base which includes the propositions that Brad will not visit and that Amy can visit (since the case asserts this), but (*per* the pragmatic rule) *excludes* the proposition that Amy will not visit, even though the case asserts this. This modal base restricts the live worlds to w_1 and w_3. Of these, w_1 is the best, and in it, Amy visits. So (1) is true at its context of evaluation. Similarly, the pragmatic rule predicts that (2) 'Brad ought to visit' is evaluated relative to a modal base which *includes* the propositions that Brad can visit and that Amy will not visit (because the case asserts this), but (*per* the pragmatic rule) *excludes* the proposition that Brad will not visit (even though the case asserts this). This modal base restricts the live worlds to w_2 and w_3. Of these, w_2 is best, and in it, Brad visits. So (2) is likewise true at its context of evaluation.

We are now in a position to see what, on the proposed Kratzerian view, accounts for the fact that (1) and (2) are both true while (3) 'Amy and Brad ought to visit' is false. On its joint-

ought interpretation, (3) says that, out of their joint option-set, [Amy visits, Brad visits] is the option which is 'to be done' by Amy and Brad together.[46] Amy and Brad *jointly* occupy the subject-position of this joint-ought claim: they *together* are its subject. In turn, this claim is evaluated relative to a modal base which *excludes* the proposition that Amy and Brad will not both visit. For the description of the case makes clear that Amy and Brad *can* both visit. So, *per* the pragmatic rule we have identified, worlds in which they both visit must be kept live in evaluating the joint-ought claim (3) 'Amy and Brad ought to visit.' Thus, unlike in (1) and (2), in the joint-ought claim (3), OUGHT quantifies over *all* of w_1-w_4: for *all* of these worlds are compatible with a modal base which includes the propositions that Amy and Brad can each visit, and can both visit; and which excludes the proposition that they will fail to visit. The (tied) best of these worlds are w_1 and w_2. In neither world do Amy and Brad both visit. So, while (1) and (2) are true relative to their contexts of evaluation, the joint-ought claim (3) is false relative to its context of evaluation.

The pragmatic rule we have identified equally explains why there is a context shift across our judgments about SLICE AND PATCH. The case has four possible outcomes, which can be ranked as follows:

BEST	**w_5** Slice slices and Patch patches
2ND BEST	**w_6** Slice doesn't slice and Patch doesn't patch
TIED WORST	**w_7** Slice slices and Patch doesn't patch
TIED WORST	**w_8** Patch patches and Slice doesn't slice

[46] *See* nn. 9-10. If to-be-doneness is understood in terms of conclusive reason, then the joint-ought interpretation of sentence (3) says that, out of their joint option-set, [Amy visits, Brad visits] is the option which Amy and Brad together have most reason to pursue.

The modal base relative to which ought-claim (5) 'Slice ought to slice' is evaluated *includes* the propositions that Patch will not patch, and that Slice can slice (since the case asserts this), but (*per* the pragmatic rule) *excludes* the proposition that Slice will not slice, even though the case asserts this. The top-ranked world consistent with this modal base is w_6, in which Slice does not slice. Similarly, (6) 'Patch ought to patch' is evaluated relative to a modal base which *includes* the propositions that Slice will not slice, and that Patch can patch (since the case asserts this) but *excludes* the proposition that Patch will not patch (*per* the pragmatic rule). Again, the top-ranked world consistent with this modal base is w_6, in which Patch does not patch. So (5) and (6) are both false at their contexts of evaluation.

However, the joint-ought claim (4) 'Slice and Patch ought to slice and patch' is true at its context of evaluation: OUGHT here quantifies over all of w_5-w_9; while, in the best of these worlds (w_5), Slice slices and Patch patches. OUGHT here quantifies over all of w_5-w_9 because – as we have argued – Slice and Patch *together* are the subject of the relevant ought-claim. The joint-ought claim (4) says that out of their joint options, [Slice slices, Patch Patches] is that which is 'to be done' by Slice and Patch. Moreover, the description of the case suggests that Slice and Patch can perform the combination of acts [Slice slices, Patch patches]. So, it is simply another application of the pragmatic rule we have identified that, in evaluating joint-ought claim (4), worlds in which Slice slices and Patch patches are treated as live.

The significance of the notion of a *joint-ought interpretation* of sentences like (4), for the proposed Kratzerian view, is now clear. On the joint-ought interpretation of (4), it expresses a joint-ought claim, whose subject is the *plurality* Slice and Patch. It is this idea – that Slice and Patch *together* occupy the subject-position of the ought-claim – which (*via* the pragmatic rule) can explain the fact that, in determining whether Slice and Patch ought to slice and patch, speakers do

not treat their failures to do so as fixed (and in turn, take sentence (4) to be true). If – *contra* our overarching proposal – the English sentence (4) 'Slice and Patch ought to slice and patch' served only to convey the pair of claims 'Slice ought to slice' and 'Patch ought to patch,' this would be puzzling. For in evaluating the former claim, we treat Patch's failure to patch as given, and in evaluating the latter, we treat Slice's failure to slice as given.

V. The Evaluative Alternative

Let us take stock. At the outset, we observed that sentence (3) 'Amy and Brad ought to visit' sounds false, even though (1) 'Amy ought to visit' and (2) 'Brad ought to visit' sound true; and that sentence (4) 'Slice and Patch ought to slice and patch' sounds true, even though (5) 'Slice ought to slice' and (6) 'Patch ought to patch' sound false. We proposed that (3) sounds false and (4) true because they admit of *joint-ought* interpretations, which are false and true respectively (§I). The joint-ought claim 'Amy and Brad ought to visit' is false because it is not the case that, of Amy and Brad's joint options, the course of action [Amy visits, Brad visits] is that which is 'to be done' by Amy and Brad.[47] The joint-ought claim 'Slice and Patch ought to slice and patch' is true, because it is the case that, of Slice and Patch's joint options, the course of action [Slice slices, Patch patches] is that which is 'to be done' by Slice and Patch (§II). We then demonstrated that both of two opposing semantic analyses of 'ought' can model the *joint-ought* interpretation of sentences like (3) and (4) (§III-IV). In closing, we turn to address a challenge to our overarching proposal, *viz.* that the best explanation of the fact that sentence (3) sounds false, and that sentence (4) sounds true, is that these sentences are being interpreted as joint-ought claims.

The challenge takes the form of a rival explanation. This rival explanation appeals to the familiar observation that 'ought' sentences admit of two sorts of interpretations: one *deliberative*

[47] *See* nn. 9, 10, and 46.

and the other *evaluative*.[48] Our focus heretofore has been the former. For we have been discussing individual-ought claims and joint-ought claims, which – as we defined those terms – pick out some course of action as *to be done* by some agent(s). According to the standard distinction, this is what 'ought' sentences do when they are interpreted in the *deliberative* sense. Interpreted in the *evaluative* sense, by contrast, an 'ought' sentence says, roughly, that it would be ideal or best if the world were a certain way.[49] Many sentences are most naturally interpreted in one or the other of the two senses.[50] For example, 'you ought to file your tax return' is most naturally interpreted in the deliberative sense, while 'there ought to be less suffering in the world' is most naturally interpreted in the evaluative sense. For the latter (unlike the former) does not say that anything is to be done by anyone. It merely says that, ideally, there would be less suffering in the world.

This distinction opens up an alternative explanation of our linguistic data. It may be proposed that the best explanation of the fact that sentence (3) sounds false, and that sentence (4) sounds true, is not that these sentences are being interpreted as joint-ought claims (and thus, as *deliberative* 'ought' sentences' of a certain kind), but instead that they are being interpreted as *evaluative* 'ought' sentences.[51] [52] Interpreted evaluatively, sentence (3) says that it would be best if Amy and Brad visited. This is not so: recall that, if Amy and Brad visit, their grandmother will become upset and confused. It is best for only one of them to visit. Thus, on its evaluative

[48] (Schroeder 2011). Others, *e.g.* (Williams 1981) and (Wedgwood 2006; 2007) , make similar distinctions.

[49] (Schroeder 2011, 1).

[50] (Schroeder 2011) argues that the *word* 'ought' is systematically ambiguous between a deliberative interpretation (on which it functions as a two-place predicate relating an agent to an action), and an evaluative interpretation (on which it functions as a propositional operator). This thesis is controversial. However, even those who reject it typically accept that whole 'ought' *sentences* admit of both deliberative and evaluative interpretations (see, *e.g.*, Finlay and Snedegar 2014).

[51] (Dietz 2016, 962–63) discusses a version of this same proposal.

[52] Note that proponents of this proposal will have to provide some explanation of why we interpret the relevant individual-ought sentences ((1), (2), (5), (6)) deliberatively, but switch to interpret sentences (3) and (4) evaluatively.

interpretation, (3) is false. Similarly, interpreted evaluatively, sentence (4) says that it would be best if Slice sliced and Patch patched. This is so: recall that if and only if they both perform their parts of the operation, Patient will survive. Thus, on its evaluative interpretation, (4) is true.

We do not deny that sentences like (3) and (4) admit of these evaluative interpretations.[53] But we *do* deny that the best explanation of our initial linguistic data – *viz*. that (3) sounds false and (4) sounds true – is that (3) and (4) are being interpreted as evaluative 'ought' sentences. The best explanation of this data, we maintain, is that (3) and (4) are being interpreted as joint-ought claims. For two reasons, our proposal is explanatorily superior.

First, it is difficult to see why the evaluative interpretations of (3) and (4) would prevail in the contexts of LONELY GRANDMOTHER and SLICE & PATCH. Although sentences with the syntactic form 'A ought to φ' *can* be read in the evaluative sense,[54] they are, as a general rule, most naturally read in the deliberative sense. And although this presumption in favor of the deliberative interpretation can be defeated, it is typically defeated by the presupposition that A is *unable* to φ. A key difference between deliberative and evaluative 'ought' sentences is that the former, but not that latter, imply 'can.'[55] The presupposition that A cannot φ renders the deliberative interpretation of 'A ought to φ' false, but has no such effect on its evaluative interpretation. Thus, on the assumption that speakers utilize a pragmatic rule of charity, it is reasonable to hold that the evaluative interpretation of 'A ought to φ' will be (less un)natural in contexts in which it is salient that A is incapable of φ-ing. Consider an example from Schroeder (2011):

[53] Note that, as (Schroeder 2011, 7–8) points out, sentences of the form 'it ought to be that A φ-s' can be syntactically rearranged into the form 'A ought to φ', by "raising" 'A' to the subject-position. We thus do not take the mere syntax of sentences (3) and (4) to establish that these sentences are interpreted deliberatively.

[54] *See* n. 53.

[55] (Williams 1981, 120–21) emphasizes this feature of the deliberative interpretation of 'ought' sentences. (Schroeder 2011, 8–11) also identifies it as one of five hallmarks of the deliberative ought, alongside: mattering directly for advice, closing deliberation, a connection with accountability, and a connection with obligation.

(9) Luckless Larry ought to win the lottery

Poor Larry has had such tough luck lately that, if anyone deserves to win, he does. But winning the lottery is not *within his power*. This fact renders the deliberative interpretation of (9) false. Since winning the lottery is not one of Larry's options – not something which Larry can do – it cannot be that action which is 'to be done' by him. Nevertheless, (9) sounds true. Larry's inability to win the lottery is salient, such that the sentence is most naturally read in its (true) evaluative sense; that is, as saying that it would be best for Larry to win.

We readily concede that certain sentences of the form 'A ought to φ,' such as (9), are most naturally read as evaluative, despite their syntactic structure. But sentences (3) and (4) are unlike these sentences in the crucial respect. Take LONELY GRANDMOTHER. Amy and Brad have the option of visiting (alone or together) – the course of action [Amy visits, Brad visits] is within their power. In turn, there is no reason to expect the deliberative interpretation of (3) 'Amy and Brad ought to visit' to be defeated. Similarly, in SLICE & PATCH, Slice and Patch have the option of slicing and patching. So, again, there is no reason to expect the deliberative interpretation of (4) 'Slice and Patch ought to slice and patch' to be defeated. In brief, because it is reasonable to expect a deliberative interpretation of 'A ought to φ' to dominate where it is presupposed that A can φ, it is reasonable to expect deliberative interpretations of (3) and (4) to dominate in LONELY GRANDMOTHER and SLICE AND PATCH. This favors our proposal, namely, that (3) sounds false and (4) sound true because they admit of *joint-ought* interpretations, which are false and true respectively. For, our proposal (but not the rival proposal) has it that – as it is reasonable to expect – sentences (3) and (4) are read deliberatively.

The second reason to favor our proposal over the evaluative alternative also invokes the relationship between deliberative 'ought' sentences and ability facts. As we have said, a hallmark

of a deliberative 'ought' sentence 'A ought to φ' is that it implies that A can φ. By contrast, an evaluative 'ought' sentence 'A ought to φ' implies nothing about A's ability to φ. Now consider a modified version of SLICE & PATCH:

> **INJURED SLICE & PATCH** Patient will die painlessly unless he receives a simple operation, which requires Dr. Slice to slice and Dr. Patch to patch. If Patient is sliced but not patched, or vice-versa, he'll die an agonizing death. Each surgeon knows that the other is going to play golf instead of doing their part of the operation. As it happens, Dr. Slice is recovering from a bad hand injury, and is unable to slice patients open.

Here, we contend, there is no temptation to say (3) 'Slice and Patch ought to slice and patch.' As soon as it is stipulated that Slice's hand is injured, such that he *cannot* slice, the intuition that 'Slice and Patch ought to slice and patch' is true evaporates. This is strong evidence that (3) is being interpreted as a deliberative 'ought' sentence. Were it being interpreted as an evaluative 'ought' sentence, its truth-value would not be sensitive to Slice and Patch's options.

In sum, because sentences (3) and (4) are most naturally interpreted as deliberative, we should, in explaining why these sentences sound false and true respectively, favor an explanation which allows that they are interpreted deliberatively. In this paper, we have provided and defended such an explanation. Sentences (3) and (4) admit of joint-ought interpretations, which are deliberative. Far from an extravagant theoretical posit, the cogency of these joint-ought interpretations look to follow from a plausible and widely accepted conception of how facts about what agents ought to do relate to facts about what their options consist in. And far from necessitating any onerous or *ad hoc* semantic engineering, joint-ought claims are readily assimilated into standard semantic models for deliberative 'ought' sentences. Joint-ought talk makes sense.[56]

[56] We are grateful to Nathan Howard, Sarah McGrath, Michael Smith, Una Stolnjic, and Christine Tiefensee for helpful discussion about the arguments in this paper, as well as to the attendees of the 2023 Princeton University dissertation

BIBLIOGRAPHY

Broome, John. 2013. *Rationality Through Reasoning*. Wiley-Blackwell.

Chaves, Rui P. 2012. "Conjunction, Cumulation and Respectively Readings." *Journal of Linguistics* 48 (2): 297–344. https://doi.org/10.1017/S0022226712000059.

Cho, Sae-Youn. 2013. "Respectively Construction in English: A Constraint- and Construction-Based Approach." *Linguistic Research* 30 (2): 273–91. https://doi.org/10.17250/KHISLI.30.2.201308.006.

Cohen, G. A. 1983. "The Structure of Proletarian Unfreedom." *Philosophy & Public Affairs* 12 (1): 3–33.

Collins, Stephanie. 2019. *Group Duties: Their Existence and Their Implications for Individuals*. Oxford University Press.

Dietz, Alexander. 2016. "What We Together Ought to Do." *Ethics* 126 (4): 955–82. https://doi.org/10.1086/686002.

Estlund, David M. 2020. *Utopophobia: On the Limits (If Any) of Political Philosophy*. Princeton, New Jersey: Princeton University Press.

Finlay, Stephen, and Justin Snedegar. 2014. "One Ought Too Many." *Philosophy and Phenomenological Research* 89 (1): 102–24.

Gawron, Jean Mark, and Andrew Kehler. 2004. "The Semantics of Respective Readings, Conjunction, and Filler-Gap Dependencies." *Linguistics and Philosophy* 27 (2): 169–207.

Gibbard, Allan. 2003. *Thinking How to Live*. Cambridge, Mass: Harvard University Press. https://www.jstor.org/stable/10.2307/j.ctv1pncqb8.

Jackson, Frank. 1985. "On the Semantics and Logic of Obligation." 1985. https://philpapers.org/rec/JACOTS.

Jackson, Frank, and J. E. J. Altham. 1988. "Understanding the Logic of Obligation." *Aristotelian Society Supplementary Volume* 62 (1): 255–83. https://doi.org/10.1093/aristoteliansupp/62.1.255.

Jackson, Frank, and Robert Pargetter. 1986. "Oughts, Options, and Actualism." *The Philosophical Review* 95 (2): 233. https://doi.org/10.2307/2185591.

Killoren, David, and Bekka Williams. 2013. "Group Agency and Overdetermination." *Ethical Theory and Moral Practice* 16 (2): 295–307. https://doi.org/10.1007/s10677-012-9336-9.

Kratzer, Angelika. 1981. "The Notional Category of Modality." In *Words, Worlds, and Contexts*, edited by Hans-Jürgen Eikmeyer and Hannes Rieser. Berlin, New York: De Gruyter.

seminar, the 2023 Chapel Hill Normativity Workshop, and the 2023 Rocky Mountain Ethics Congress. Special thanks to Alisabeth Ayars, Gideon Rosen, and Mark van Roojen, who provided detailed comments on earlier drafts.

———. 2012. “Chapter 2 The Notional Category of Modality.” In *Modals and Conditionals: New and Revised Perspectives*, edited by Angelika Kratzer, 0. Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199234684.003.0002.

Linnebo, Øystein. 2022. “Plural Quantification.” In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Spring 2022. Metaphysics Research Lab, Stanford University. https://plato.stanford.edu/archives/spr2022/entries/plural-quant/.

McKay, Thomas. 2006. *Plural Predication*. Oxford, New York: Oxford University Press.

Mellor, Rowan. Forthcoming. "Joint Ought" *Philosophy & Public Affairs*.

Oliver, Alex, and Timothy Smiley. 2016. *Plural Logic: Second Edition, Revised and Enlarged*. Oxford University Press.

Parfit, Derek. 1988. “What We Together Do.”

Pinkert, Felix. 2014. “What We Together Can (Be Required to) Do.” *Midwest Studies In Philosophy* 38 (1): 187–202. https://doi.org/10.1111/misp.12023.

Portmore, Douglas W. 2011. *Commonsense Consequentialism: Wherein Morality Meets Rationality*. Oxford University Press USA.

———. 2019. *Opting for the Best: Oughts and Options*. Oxford Moral Theory. Oxford, New York: Oxford University Press.

Ralph Wedgwood. 2007. *The Nature of Normativity*. Oxford: Oxford University Press.

Rayo, Agustín. 2007. “Plurals.” *Philosophy Compass* 2 (3): 411–27. https://doi.org/10.1111/j.1747-9991.2007.00060.x.

Schroeder, Mark. 2011. “*Ought* , Agents, and Actions.” *The Philosophical Review* 120 (1): 1–41. https://doi.org/10.1215/00318108-2010-017.

Schwenkenbecher, Anne. 2019. “Collective Moral Obligations: ‘We-Reasoning’ and the Perspective of the Deliberating Agent.” *The Monist* 102 (2): 151–71. https://doi.org/10.1093/monist/onz003.

———. 2021. *Getting Our Act Together: A Theory of Collective Moral Obligations*. New York: Routledge.

Shea, Margaret. forthcoming. “Why Plan-Expressivists Can’t Pick Up the Moral Slack.” *Oxford Studies in Metaethics* 19.

Smith, Thomas H. 2009. “III-Non-Distributive Blameworthiness.” *Proceedings of the Aristotelian Society (Hardback)* 109 (1pt1): 31–60. https://doi.org/10.1111/j.1467-9264.2009.00257.x.

Spiekermann, Kai. 2021. “Group Duties: Their Existence and Their Implications for Individuals, by Stephanie Collins.” *Mind* 130 (518): 714–23. https://doi.org/10.1093/mind/fzaa004.

Wedgwood, Ralph. 2006. “The Meaning of ‘Ought.’” 2006. https://philpapers.org/rec/WEDTMO-3.

White, Stephen (forthcoming). Collective Obligations and the Ethics of Participation. In: Kyla Ebels-Duggan & Berislav Marúsiç (eds.) *Responsibility and the Demands of Morality*. Oxford University Press.

Williams, Bernard, ed. 1973. “Ethical Consistency.” In *Problems of the Self: Philosophical Papers 1956–1972*, 166–86. Cambridge: Cambridge University Press. https://doi.org/10.1017/CBO9780511621253.013.

———, ed. 1981. “Ought and Moral Obligation.” In *Moral Luck: Philosophical Papers 1973–1980*, 114–23. Cambridge: Cambridge University Press. https://doi.org/10.1017/CBO9781139165860.010.

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