

CHAPTER 11

UNDERSTANDING WHAT WE OUGHT AND SHALL DO A HYPERSTATE SEMANTICS FOR DESCRIPTIVE, PRESCRIPTIVE, AND INTENTIONAL SENTENCES

Forthcoming in *Groups, Norms, and Practices*, edited by Ladislav Koreň, Hans Bernhard Schmid, Preston Stovall, and Leo Townsend (Cham: Springer).

Preston Stovall

Philosophy Faculty, University of Hradec Králové

preston.stovall@uhk.cz

Abstract: This essay is part of a larger project aimed at making sense of rational thought and agency as part of the natural world. It provides a semantic framework for thinking about the contents of: 1) descriptive thoughts and sentences having a representational or mind-to-world direction of fit, and which manifest our capacity for theoretical rationality; and 2) prescriptive and intentional sentences having an expressive or world-to-mind direction of fit, and which manifest our capacity for practical rationality. I use a modified version of Allan Gibbard's *hyperstate semantics*, employing both maximally determinate possible worlds and maximally determinate plans of action, as a basis for providing a unified understanding of moral judgments and expressions of individual and collective intentionality – they one and all give voice to our ability as rational agents to adopt the perspectives of various individuals within a community and consider how we would behave were we in their positions. In the course of spelling out the view I draw on and criticize ideas that Wilfrid Sellars advanced in the middle of the 20th century, while employing the tools of contemporary modal logic and model-theoretic semantics to give perspicuous formulation to his thought that the moral judgment 'one ought to do A' should be understood in terms of the collective intention 'we shall do A', where the pronoun denotes the unrestricted class of rational agents and the 'ought' is in some sense unconditionally binding.

Keywords: collective intentionality, deontic modality, expressivism, planning semantics, Sellarsian ethics

11.1 INTRODUCTION: THE MORAL 'OUGHT' AND THE INTENTIONAL 'SHALL'

Though he is not often discussed alongside the 'big four' of collective intentionality (Bratman, Gilbert, Searle, and Tuomela)¹, Wilfrid Sellars did much to stimulate interest in the subject in the middle of the 20th century. Over the course of a series of works addressing the nature of moral reasoning, Sellars developed an account of collective intentionality on the

¹ For the reference see the Introduction to Chant, Hindriks, and Preyer 2014. For exceptions see Koons (**this volume**) Olen and Turner (2015), Peregrin (**this volume**), and Schweikard and Schmid (2013) §2.3.

basis of which to understand moral judgment (e.g. 1951, 1963, 1966a, 1966b, 1967a, chapter 7 of 1968/1992, 1976, and 1980). In doing so he outlined (here one could not say *developed*) a theory of the semantic functions of modal terms like ‘shall’, ‘ought’, and ‘could’. His idea is that the assertion of a sentence like ‘bankers ought to dress professionally’ gives voice to a collective intention concerning what anyone in the community of bankers *shall* do. Moral judgments are or involve a universalization of this collective intentionality across all rational beings: use of the moral ‘ought’ is a way of giving voice to what any of us shall do where the ‘us’ is unrestricted and where ‘shall’ connotes an intention-in-expression (as opposed to a claim that such an intention exists). This idea relates judgments about what is right and wrong concerning exercises of instrumental practical rationality – which concern what anyone ought to do to achieve a particular end in a given situation – to judgments about what is right and wrong concerning morality. And while the instrumental ‘ought’ concerns action that is directed at some end, and so its universality is circumscribed by that end as a condition on the force of the ‘ought’, the moral ‘ought’ is one whose force is in some sense unconditionally binding.

In this regard Sellars sees himself spelling out a conception of human cognition consonant with Kant’s.² Looking back, we can also understand this as an expressivist or non-cognitivist analysis of moral judgment.³ While I will argue below that we have good reason to dispute some of what Sellars concludes about moral reasoning, collective intentionality, and the logic of the associated modals, I believe that in some areas he saw farther and more clearly than is commonly appreciated. The last chapter of *Science and Metaphysics* contains a discussion of the different roles played by the *contents* of collective intentions as against their

² From the closing chapter of *Science and Metaphysics* (1968), a book whose subtitle is “Variations on Kantian Themes” (p.176):

...in my opinion the metaphysics of morals is but a fragment of a broader critique of practical reason.

³ Though the term ‘expressivism’ was not then in common currency, Sellars occasionally discusses non-cognitivism (usually in connection with emotivism) – cf (1967b). See Klemick (2018) for a quasi-realist defense of Sellars’ non-cognitivism about ethics, and Koons (2019), chapter 5, for a general discussion of Sellars’ expressivism.

mode as collective, and of the two complementary directions of fit that characterize descriptive and intentional states of mind, that resonates today (its echoes of a classical pragmatist use of the scientific categories employed to understand neural activity as a basis for thinking about theoretical and practical rationality, however, have gone almost totally silent in the intervening years). From pp.188-9:

33....I can intend that someone else do A (intend him to do A), and he can intend to do A....There is, nevertheless, an asymmetry. For, obviously, only *his* intention can grow directly into *his* volition and *his* action. My intention with respect to his action can grow into a volition only if practical reasoning draws a conclusion from it concerning some influence I can bring to bear.

34. This latter consideration highlights the obvious fact that even where the descriptive content of two people's intentions is in the strongest sense the same not only are the intendings numerically different, which is true even in the case of belief, but the total content of the intendings involves a special mode of egocentricity (expressed by the word 'shall') which is, in practical discourse, the counterpart of the egocentricity of demonstratives. The latter is the egocentricity involved in the impact of the world on discourse, the former is the egocentricity involved in the impact of discourse on the world in volition.

With the benefit of hindsight, however, and looking across decades of work on model theory as a basis for semantically evaluating different vocabularies, we can see not only that Sellars made some conceptual moves that are not forced upon us, but also that his survey of the terrain needs to be re-examined. For Sellars left at best an outline for a semantics of moral

and intentional discourse, and while Koons (2019) offers a systematic elaboration and defense of Sellars' ethics, and relates Sellars' ethics to the work of Bratman, Gilbert, Tomasello, and Tuomela on collective intentionality, we still lack a semantic framework sufficient to unify the prescriptive and intentional modalities under a common conception of practical rationality (I use 'deontic' and 'prescriptive' interchangeably in this essay, concerning what agents ought, may, and are forbidden from doing). There are also problems with the sketch of an account Sellars left. By simply identifying deontic judgments with expressions of collective intentions (e.g. 1963, part 12, and 1968, p.219; Koons, chapter 5, also identifies 'ought' judgments with 'shall' judgments), it would appear that 'we ought to A and we shall not A' is rendered as the contradiction 'we shall A and we shall not A'. But though an assertion of the first sentence may involve some sort of practical irrationality, there is a non-empty mental state associated with it, as we've recognized for millennia – judgments such as these led the prophets to rend their clothes at what their communities were doing. Where negation is a complementarity operator, however, an assertion of the second sentence has no associated mental state (I am glossing over issues of internal and external negation that are discussed in detail in sections **11.3.2** and **11.3.3**). At the same time, Sellars thinks that a judgment that one *shall* do something commits one to the claim that one *will* (e.g. at section 4 of 1963; though rejected by Ludwig 2016, p.115, this view is not uncommon in the literature on collective intentionality – cf. Bratman 2014 p.76, Mele 1989 p. 19 and the references in note 1, and Tuomela and Miller 1988 p.375).⁴ If we allow this implication to hold, and we identify 'shall' with 'ought', we need to make sure that commitment to the claim that one *ought* to do something does not also commit one to the claim that one *will* – for just as surely as 'we

⁴ Bratman (2014) notes that this condition is too strong to be a necessary condition on intention, Mele (1989) defends the weaker claim that if one intends to do something one at least does not believe that one will not, Tuomela and Miller (1988) hedge with a parenthetical claim that one at least believes that one *probably* will do what one intends, and in footnote 23 of (1963) Sellars includes a *ceteris paribus* clause. These weaker commitments could be represented here instead of the stronger one, but as all are optional I stick to the simpler formulation so as to illustrate how to accommodate the view for those so inclined.

ought to do A and we shall not' is coherent, so is 'we ought to A and we will not'. This again suggests distinguishing 'ought' from 'shall'. Additionally, if deontic judgments simply express collective intentions it is not clear how to understand the claim that something is *permitted* but not obliged (cf. Koons 2019 pp.132ff.). And what sort of intentional state of mind is expressed when one says that the performance of some general act is obligatory, two of its specific instances are separately permitted, while one intends to do one or the other – e.g. 'I ought to go to the lecture, I am permitted to sit in the front, I am permitted to sit in the back, and I shall sit in the front'?

This essay grounds a unified conception of moral reasoning, individual and collective intentionality, and practical rationality on a solid semantic footing by using the resources of model theory and modal logic to supply an interpretation for a language that includes descriptive vocabulary, prescriptive vocabulary, and intentional vocabulary. On this semantics representational interpretations are provided for descriptive vocabulary and expressivist interpretations are provided for prescriptive and intentional vocabulary. Using a modified version of the plan-theoretic semantics presented in Gibbard (2003), my goal is to explicate the idea that both moral reasoning and collective intentionality can be understood as the exercise of practical rationality across perspectives of person, place, and time. The result is a view on which descriptive sentences represent the space of possible worlds, while prescriptive and intentional sentences express commitment to plans of action that involve putting oneself into the shoes of others within a community. The view diverges from Gibbard and Sellars in various ways (some of which I detail below) but the upshot is a semantics able to establish a number of implication relations that the descriptive 'will', the prescriptive 'ought', and the intentional 'shall' ought to conform to. I believe this framework will prove useful in sorting out some of the debate in the literature on collective intentionality, but my aim in this essay is to present the semantics, explain the resources its models employ, discuss

Sellars' account alongside my own, and motivate the conclusion that there is more to be learned from studying views of this sort. With the exception of a brief discussion of Bratman's views on planning mental states as a basis of collective intentionality in section **11.3.8**, I defer to future work a comparison of this semantics with the views of the big four and the surrounding literature.

Two final programmatic notes. First, my focus here lies solely on intentions for action, and I leave questions about collective attention, thought, and other mental states undiscussed. Second, as my interest lies in the modal reasoning and judgment that attends practical rationality I do not discuss modalities for non-agentive normative evaluation, e.g. concerning how organisms and artifacts ought to be or behave.

11.2 A DEONTIC HYPERSTATE SEMANTICS

11.2.1 BACKGROUND AND FRAMEWORK

In other work (Manuscript) I show that by taking a class of atomic descriptive sentences interpreted via possible worlds, a class of atomic prescriptive sentences interpreted via plans of action of the sort discussed below, and the standard set-theoretic definitions for the Boolean operators, the content of any mixed or pure descriptive and prescriptive sentence can be recursively modelled, where the mind-to-world descriptive function and the world-to-mind prescriptive function is kept distinct. In this part of the essay I present an abbreviated formulation of that semantics, and in **11.3** I discuss its extension to account for intentional sentences.

The central semantic concept is the notion of a *hyperstate*, introduced in Gibbard (2003) and understood as an ordered pair of a maximally determinate possible world and a maximally determinate plan of action (a *hyperplan*). Think of a hyperplan as a specification of what a person would do at every possible point of choice, including situations where one is

other people – and so a hyperplan specifies what to do if one were Caesar faced with crossing the Rubicon, or a colleague faced with the choice of whether or not to attend a lecture. Just as descriptive mental states, having a mind-to-world direction of fit, can be understood as sets of possible worlds (the worlds at which the descriptions are true), so can prescriptive mental states, having a world-to-mind direction of fit, be understood as sets of plans of action (the plans where the agent, or the members of a class of agents, chooses to do the thing prescribed). Because a hyperstate includes both a world and a plan, it is possible to distinguish various features of the contributions that descriptive and prescriptive or deontic thoughts make to our cognitive economy.

While the semantic values of descriptive and prescriptive sentences can be provided on this basis, a straightforward extension of this framework would use these hyperstates to model the contents of intentional sentences as well. This leads to the problem that the sentence ‘we ought to A and we shall not A’ becomes the contradiction ‘we shall A and we shall not A’, however. And so in **11.3** I distinguish deontic hyperplans from intentional hyperplans, and I define a *deontic-intentional hyperstate* as an ordered triple of a possible world, a deontic hyperplan, and an intentional hyperplan. But I begin here because the ordered pairs – and the descriptive/prescriptive fragment of the language – are simpler to work with.

11.2.2 SINGLE-MINDEDNESS AND INDIFFERENCE

The deontic hyperplans that model prescriptive mental states require that we distinguish two different ways of choosing – single-mindedly and indifferently. To forecast, one chooses an action single-mindedly when one regards it as obliged and one chooses an action indifferently when one regards it as merely permitted – thus, belief that one ought to attend the lecture and that one is permitted to sit in either the front or the back of the room will be understood as a

plan to choose single-mindedly to attend the lecture while choosing indifferently whether to sit in the front or the back. The distinction between these two choice attitudes has gone neglected in the literature on planning expressivism. While it is discussed in Gibbard (2003), that discussion has been seldom mentioned by his critics. Silk (2015 p.53, footnote 8) notes that by passing over this distinction many of Gibbard's critics have not done service to his view. Dreier (2006, 2009) advances a similar distinction, using the same terminology, but without noting its occurrence in Gibbard. Interestingly, Sellars introduced the same distinction in his own work (e.g. 1966b pp.113-115 and 1976 p.57), using the very same terms that both Gibbard and Dreier use, though to my knowledge he never discusses it in any detail (after introducing the distinction in the latter essay Sellars writes "but I am concerned to sketch rather than map"). As neither Gibbard nor Dreier note that Sellars discusses this distinction, it appears that the notions of single-minded and indifferent choice attitudes have been independently developed at least three times in the last half century.⁵ These notions are absolutely central to a planning semantics for the deontic modalities, however. For in order to distinguish the strong and weak deontic modalities of obligation/forbiddance and permission, each point of choice on a deontic hyperplan must be specified according to whether the decisions undertaken are done so either single-mindedly or indifferently. The hyperplans that model intentional mental states do not need to mark this distinction, however (why this is so will become clear below). That is, each point of choice on an intentional hyperplan (introduced in **11.3**) specifies only which actions are taken and which are not, without regard for the attitudes under which the choices are made. This indicates, I will suggest, that the intentional state of mind, whether individual or collective, is a less cognitively demanding and conceptually sophisticated sort of mental state than the state we express when we issue

⁵ Sellars, Gibbard, and Dreier each speak of *preferential* and indifferent choices, but because one might suppose that any choice exhibits one's preferences I choose to contrast indifference with single-mindedness. And as I provide my own definition of indifference in terms of single-mindedness I will use this terminology.

prescriptions under the deontic modalities (this also motivates keeping two sets of books for deontic and intentional plans).⁶

The notion of a single-minded choice is to be understood through the attitude of rejection, and the indifferent choice can be introduced on this basis (Gibbard introduces talk of rejection, but the definition of indifference in terms of rejection and single-mindedness is my own). Just as one can *believe* that some state of affairs obtains, so can one *reject* some course of action. An agent chooses an action A single-mindedly when she rejects every action B that is incompatible with A (rejection is a term of art in this sense). An agent chooses an action A indifferently just in case there is some action B that is 1) incompatible with A and 2) is such that she could have picked B without changing any of her single-minded choices. When an agent makes a choice indifferently she has *rejected rejecting* each of the choices that she could take without changing any of her single-minded choices. Thus, it is of the nature of both single-mindedness and indifference that one who chooses from these frames of mind is at the same time attitudinally related to choices she does *not* undertake. The intentional choice may, however, be specified solely in terms of the choice itself; this again suggests that the deontic frame of mind involves a sort of cognitive sophistication that the merely intentional does not.

It is clear that we can be in these planning states of mind, and the mental states associated with the assertions of sentences concerning what is forbidden, what is obliged, and what is permitted can be specified in terms of them. For instance, suppose I have decided that I ought not go home early this afternoon, that I ought to go to the department colloquium, and that it is both permissible that I sit in the front of the room and permissible that I sit in the back of the room. In this circumstance my mental state includes rejecting going home and rejecting not going to the talk, modelled by the set of hyperplans where I single-mindedly

⁶ I examine relations of conceptual and ontological priority among deontic, individually intentional, and collectively intentional mental states in (Forthcoming).

choose to go to the talk by rejecting every choice that is incompatible with going to the talk – including, among other things, the choice to go home. If I do not change my mind when it comes time to attend the talk, then I actualize this planning mental state by choosing single-mindedly to go to the talk. Finally, believing that both sitting in the front of the room and sitting in the back of the room are permitted is a mental state modelled by the set of plans where I variously choose indifferently to sit in the front of the room and indifferently to sit in the back of the room. When it comes time to act on one plan rather than the other I actualize the plan in question by choosing indifferently between those choices, which is to say I *reject* rejecting the alternative (the option I could have undertaken without changing any of my single-minded choices). And so should I choose to sit in the front of the room that choice is both an indifferent choice to sit in the front of the room and an indifferent choice not to sit in the back of the room.

11.2.3 A HYPERSTATE SEMANTICS FOR DESCRIPTIVE AND PRESCRIPTIVE SENTENCES

In order to model sentences concerning what restricted classes of people ought to do in specific circumstances, it will occasionally be useful to talk about *types* of circumstances (weeknights, weeknights where there is a colloquium talk, etc.), *kinds* of actions (walking to a colloquium talk, walking quickly to a colloquium talk, etc.), and *classes* of agents (all of the people in my department, all philosophers, etc.). Tradeoffs can be made in terms of where we locate complexity, and for the semantics of prescriptive sentences I will assume that circumstances include information about the classes of agent(s) who make choices there, while keeping kinds of action separate from the circumstances in which they occur.⁷ E.g., deontic hyperplans concerning how everyone in my department should dress for work will be

⁷ Sellars (1968, p.195) treats hypothetical imperatives in much the same way, holding that they apply to all rational beings insofar as they find themselves in certain circumstances (see the discussion at Koons 2019 pp.72ff.).

framed in terms of what the unrestricted class of agents would do were they to find themselves in a circumstance where they are members of my department faced with the choice of how to dress for work. A language that included kind-restricted quantifiers could be supplied a semantics that distinguishes the non-universal classes of agents to which prescriptions apply from their circumstances of action. The introduction of quantificational devices would also facilitate a treatment of *de re* and *de dicto* modality, but doing so would require relativizing semantic values to points of evaluation and such a discussion would take us too far afield (see the **Appendix** for a brief overview of these issues). In future work I hope to examine the interaction between quantification and modality on a hyperstate semantics in more detail, but here I treat prescriptions as applying universally and build their restricted applications into the circumstance condition. Additionally, all deontic modality over sentences concerning groups of people should be read *de dicto* (though again, see the **Appendix** for a discussion of *de re* and *de dicto* modality).

Consider a language \mathbf{L} consisting of atomic descriptive sentences, atomic sentences falling under one of the deontic modalities of obligation, permission, and forbiddance, and the sentences composed from these together with the Boolean operators. Hyperstates can be used to provide a semantic evaluation on which the descriptive sentences represent the space of possible worlds and the prescriptive sentences express plans of action. Let a *possible world* w be defined as a maximally consistent set of true sentences. That is, for every descriptive sentence $d \in \mathbf{L}$ and every w , either d is true at w or $\sim d$ is true at w . The content of a descriptive sentence can then be understood as the set of worlds at which it is true. In this regard descriptive sentences can be understood as representations of the space of possible worlds. Plans of action are defined in terms of the choices made on them, where these plans are maximal in the sense that they specify what anyone would do at any every of choice. More precisely, let a *deontic hyperplan* h_D be defined as a maximally consistent plan of action

such that, for every circumstance C , every agent α able to make a choice at C , and every action A that α is able to choose to perform at C , either (exclusively):

- 1) α single-mindedly chooses to A at C on h_D ; or
- 2) α single-mindedly chooses not to A at C on h_D ; or
- 3) α indifferently chooses to A at C on h_D ; or
- 4) α indifferently chooses not to A at C on h_D .

Officially α is a metalinguistic variable but I abuse notation and occasionally treat it as an instance in the object language (e.g. in **11.3.8**).

Various requirements that a semantics must meet in order to model our notion of agency correspond to restrictions placed on the relations among hyperplans. For instance, notice that in general one undertakes many actions at a single point of choice—viz., to choose to attend the department colloquium this afternoon is also to choose *not* to go home early. Whatever determinate inclusions and exclusions may obtain among choices should be reflected in the plans we use to think about them. It would also seem part of the logic of agency that wherever it is possible to plan to do something then it is also possible to plan not to do it, understood *de dicto* in the sense that one could plan to do what one takes to be something else. In the branching-times framework of Belnap, Perloff, and Xu (2001), employing the ‘sees to it that’ modal operator for talk of agency, this is represented by a requirement that where it is possible to see to it that something is the case it is also possible to not see to it that it is the case – there is no agency where an agent could not have done otherwise (this is the ‘negative requirement’ on agency, discussed at pp.36-7). If that is right, and if our talk of intention inherits the logic of our talk of agency, then the following is a plausible condition to impose on the notion of a hyperplan. With ‘ h_1 ’ and ‘ h_2 ’ ranging over

deontic and intentional hyperplans (introduced in **11.3**), and with ‘chooses*’ ranging over ‘chooses indifferently’, ‘chooses single-mindedly’, and ‘chooses’ (the third case is to ensure that intentional hyperplans also meet this condition):

GENUINE AGENCY: $\forall \alpha \forall A \forall C \forall h_1$ [if α chooses* to A at C on h_1 then $\exists h_2$ (α chooses* not to A at C on h_2)]

GENUINE AGENCY says that for any plan on which one chooses to do something (perhaps under a particular attitude) there is another plan on which one chooses (under that same attitude) not to do so.

Let a *deontic hyperstate* $\langle w, h_D \rangle$ be an ordered pair of a possible world and a deontic hyperplan and let S_D denote the set of deontic hyperstates. That is, where ‘W’ denotes the set of possible worlds and ‘ H_D ’ denotes the set of deontic hyperplans, a deontic hyperstate $\langle w, h_D \rangle \in S_D$ is an element of the Cartesian product of W and H_D . The contents of states of mind voiced with descriptive and prescriptive sentences are given by the world and choice components of deontic hyperstates (from here on I tend to trade talk of states of mind for the contents of the sentences used to convey them). Where $[[\phi]]$ denotes the content of a sentence $\phi \in \mathbf{L}$ under a given interpretation, for descriptive sentences d we have:

$$[[d]] =_{\text{def.}} \{ \langle w, h_D \rangle : d \text{ is true at } w \}$$

For prescriptive sentences we distinguish three varieties, corresponding to deontic judgments concerning what is forbidden, what is obliged, and what is permitted. Schematically, the contents of the three kinds of deontic sentence are determined as follows. An assertion of ‘people are obliged to A in C’ expresses universally rejecting not doing A in C:

[[people are obliged to A in C]] =_{def.} {<w, h_D>: for every α , α single-mindedly chooses to A in C on h_D}

Thus, the content of ‘my colleagues are obliged to go to colloquia’ is the set of hyperstates whose hyperplans are such that every person who is in a circumstance of being my colleague and faced with a choice of either going or not going to a colloquium chooses single-mindedly to go. This will include, e.g., plans where a colleague chooses indifferently to sit in the front of the room, chooses indifferently to sit in the back, etc., so long as these choices are also ones where the agent chooses single-mindedly to go to the colloquium. To say that my colleagues are obliged to go to colloquia, then, is to express commitment to a plan concerning what I would single-mindedly do were I anyone of my colleagues.

Similarly, an utterance of ‘people are forbidden to A in C’ expresses universally rejecting doing A in C:

[[people are forbidden to A in C]] =_{def.} {<w, h_D>: for every α , α single-mindedly chooses not to A in C on h_D}

And so the content of the thought given voice with an utterance of ‘my colleagues are forbidden from going to colloquia’ is the set of hyperstates whose hyperplans are such that every person who is my colleague and faced with the choice of going to a colloquium chooses single-mindedly not to go.

Finally, the assertion of ‘people are permitted to A in C’ expresses universally rejecting rejecting doing A in C:

[[people are permitted to A in C]] =_{def.} {<w, h_D>: for every α , α either single-mindedly chooses to A in C or indifferently chooses whether or not to A in C on h_D}

The content of the thought that my colleagues are permitted to go to colloquia is therefore the set of hyperstates whose hyperplans are such that every colleague of mine who is in a circumstance of being faced with attending a colloquium *either* chooses single-mindedly to go, *or* chooses indifferently whether or not to go (this allows for individual hyperplans where an agent sometimes chooses to indifferently to go and sometimes chooses indifferently not to go). In this case the only hyperstates ruled out are those whose hyperplans include colleagues who choose single-mindedly not to go. It is easily established that where negation is a complementarity operator, understood in terms of rejection for prescriptive sentences, these three kinds of state of mind bear all the right duals with negation (when something is not forbidden it is permitted, when it is not obliged one is permitted not to do it, etc. – see Manuscript for details).

11.3 A DEONTIC-INTENTIONAL HYPERSTATE SEMANTICS

11.3.1 FROM DEONTIC PRESCRIPTION TO ‘WE-SHALL’ INTENTION

According to this proposal, when everyone in my department shares the mental state of being in agreement about how we ought to comport ourselves at a colloquium, we are in a state of mind wherein we each agree on what we would do were we anyone else in the department.

To plan in this sort of way is to think of oneself as a member of a group whose other members we relate to *as if* we could make their choices for them. This feature of the planning apparatus discharges a key commitment in Sellars’ theory of moral judgment. For on this account a moral judgment that one ought to do A expresses a plan, were one *anyone*, *anywhere*, *at any time* (where A-ing is at issue) to choose to do A. And so the capacity for

reasoning in a moral frame of mind is the universalization, over person, place, and time, of the practically rational capacity to plan one's life. In this section I extend this treatment by providing a hyperstate semantics for individual and collective intentions.

11.3.2 SELLARSIAN 'SHALL'S

As I mentioned at the start of the essay, I think we have good reason to question some of the commitments Sellars takes on. Perhaps most surprisingly, he held that judgments using the 'shall' operator do not fall under the scope of logical operators. He also repeatedly examines problems with negation without arriving at a satisfactory picture of the interaction between negation and intentional vocabulary. Sellars' reasoning for denying that 'shall' sentences logically compose with other sentences is not above criticism. For one thing, it raises problems of compositionality for a language containing both descriptive and non-descriptive sentences. Koons (2019, pp.132ff.) examines this point in some detail and concludes that we should allow for embedded 'shall' sentences. It is also not clear that Sellars, by his own lights, should bar 'shall' from embedding. For he is happy to talk about *implications* among statements of intention. In a footnote at p.115 of (1966b) he defends this by asserting that "the *relation word* 'implies' ...is not a truth-functional *connective* [and] is in the metalanguage, and *mentions* rather than *uses* the shall-sentences" (emphasis in the original). This is a curious position for one who, thirteen years earlier (1953), had argued that metalinguistic claims about inferential relations among sentences could be understood in terms of object-language claims that use conditionals. And so insofar as Sellars allows metalinguistic implications like

'I shall walk quickly to the colloquium talk' implies 'I shall walk to the colloquium talk'

one would expect him to allow object-language sentences like

if I shall walk quickly to the colloquium talk, then I shall walk to the colloquium talk

Perhaps that sounds stilted. But there would seem to be a perfectly sensible state of mind associated with that conditional. At any rate, I mean to allow ‘shall’ sentences to embed arbitrarily under the Boolean operators and I defer a consideration of non-Boolean implication for another time.

There are also issues with Sellars’ discussion of negation, as it was his contention that ‘shall’ sentences cannot be externally negated in any sense that differs from their internally-negated variants. Sellars’ discussion of external and internal negation is more substantial, however, and he examines the issue across a number of texts (especially 1963, 1966b, and chapter 7 of 1968; see also Klemick 2018 §2 and Koons 2019, chapter 5, the latter containing an extended discussion of Sellars’ standpoint on negation and ‘shall’). His central concern is that ‘ought’ clearly distinguishes an external from an internal negation: ‘it is not the case that you ought to A’ differs in meaning from ‘you ought not A’ – the former marks a permission, the latter something forbidden. This is a problem for Sellars because he wants to analyze ‘ought’ in terms of the collective ‘shall’, and he does not think that ‘I shall not A’ and ‘it is not the case that I shall A’ marks a similar sort of distinction. His commitment to the equivalence between ‘ought’ and ‘shall’ thereby requires that he explain why, as he claims, expressions of individual intention do not distinguish internal and external negation while expressions for collective intentions, or the deontic expressions that go proxy for them, do.⁸

In parts 11 and 12 of (1963), an essay that is cited at a footnote of p.185 of (1968) as the point

⁸ There are other features of Sellars’ account that a fuller story would have to look out. For instance, I will have nothing to say about Sellars’ claim that intentions can at most conflict but not contradict, nor his understand of the relationship between truth and negation.

of contact for his understanding of negation in practical contexts, Sellars argues that ‘ought’ has an external negation because of the intersubjectivity of ‘ought’ claims. At the end of part 12 he writes (emphasis in the original; notice the use of ‘ordinary’ in the second sentence, apparently referencing individual rather than collective intentions – my thanks to Jim O’Shea for pointing this out):

Let me now bring all these considerations together. I suggest that the fact that ought-statements, unlike ordinary shall-statements, have a proper negation is built on the shared intending expressed by “ought”. In other words, the *syntactical* intersubjectivity of ought-statements which [makes it possible that two ought-statements contradict each other], and which *consists in* the existence of the form “not-ought[X doing A]” in addition to the form “ought[X *not* doing A]” rests on the intersubjectivity of the intention expressed by ought-statements.

It is true that expressions of collective intention, as with expressions of moral judgment, share the feature that they are in principle part of an intersubjective space of practical reasoning within some community. But we can register this point without conceiving of moral reasoning as a species of collective intentionality, and when we do we free ourselves of the need to explain why ‘ought’ has an external negation while ‘shall’ does not. Sellars faces that need only because he defined ‘ought’ in terms of ‘shall’—in the paragraph preceding the one just quoted Sellars writes (emphasis in the original) “*I suggest that ought, as an expression of intention, is a special case of shall_{we}*”. Koons (2019), chapter 5, follows Sellars in this identification and argues that we should distinguish external from internal negation for ‘shall_{we}’. I do not think we should identify ‘ought’ with ‘shall_{we}’, or otherwise define the former in terms of the latter. These two modes of expression can be

understood in terms of the general notion of planning, and in that regard they are species of the genus *practical rationality*. But as species they are differentiated separately – whereas the plans that model the deontic modalities distinguish the single-minded from the indifferent choice, the plans that model the intentional modalities do not.

11.3.3 ISSUES OF ‘SHALL’ AND ‘NOT’

As I see it, there are three questions in play here. First, there’s the question of whether we can frame a semantics on which sentences containing ‘shall’ and ‘ought’ have the right kinds of logical properties. This is a question to be decided by formal considerations and the metaphysics of the models we use to interpret the language under consideration, together with whatever arguments we give for whatever logical properties we think the semantics ought to have. Second, there is the question of whether ‘shall’ has a dual, as ‘obligation’ does with ‘permission’ and ‘necessity’ does with ‘possibility’. An answer to this question will be closely connected to the answer one gives to the first, on account of the fact that internally and externally negated ‘ought’ claims clearly differ in meaning, where the strong force of forbiddance, defined by an internally-negated ‘ought’ claim (one ought not A—i.e., it is forbidden to A) is distinguished from the externally-negated ‘ought’ claims that mark the weak force of permission (it is not the case that one ought to A—it is permitted not to A). Finally, there is the question of whether saying that one shall not do something is equivalent to saying that it is not the case that one shall do it – this is a question of whether ‘shall’, in either its individual or collective guise, as a matter of the grammar of English (or the pidgin-English spoken by a community of philosophers), distinguishes an internal and an external negation.

Concerning the first question, the GENUINE AGENCY condition and the fact that intentional hyperplans do not distinguish single-minded from indifferent choices ensures that

for any action A an agent is capable of undertaking, an externally-negated ‘shall’ sentence (whether individual or collective) is equivalent to an internally-negated one (see proof in **11.3.7**).⁹ And so from plausible assumptions about the nature of our planning agency it follows that there is no distinct external negation for ‘shall’. While this entailment is grounded in the metaphysics of agency, there is no corresponding entailment that the absence of an intention to do something is the intention not to do it, however. That second entailment corresponds to a descriptive claim about the ontology of intentions, and it would be reflected in a relation among the world elements of hyperstates. But the equivalence between ‘shall not’ and ‘not shall’ is a claim about the logic of sentences *expressing* intentional agency: any time one expresses that it is not the case that one shall do A, one has thereby expressed that one shall not do A. Nevertheless, we can preserve a distinct internal and external negation for the deontic modalities, and this feature of the language is captured by the fact that the hyperplans used for interpreting the expressions of deontic judgment specify whether a choice is made single-mindedly or indifferently. The distinction between single-mindedness and indifference thereby answers the second question by ensuring that the duality of the deontic modalities is semantically marked, while the absence of that distinction for the semantic evaluation of the intentional operator ‘shall’, together with plausible assumptions about the nature of agency, ensures that ‘shall’ has no dual.

It is of note that the absence of a discussion of Gibbard’s use of this distinction between two ways of choosing led Schroeder (2008) and Dreier (2006, 2009) to accuse Gibbard (2003) of being unable to distinguish internal and external negation for the deontic modalities. But that distinction in choice types is needed for the deontic modalities precisely so as to account for the difference between these two sorts of negations, and the associated

⁹ The restriction to actions the agent is capable of undertaking rules out cases like ‘it is not the case that I shall make the sun explode’, which should be understood as descriptive claims that one lacks a particular intention (the expressive intentional thought can be captured with ‘it is not the case that I shall *try*, etc.). If we allow hyperplans to include actions that are apparently impossible, this restriction can be lifted.

difference between the strong and weak deontic modality. But where there is no corresponding difference in negation (and no associated duality) for the intentional operator 'shall', there is no need to distinguish single-mindedness from indifference in mere intention, whether individual or collective. By not properly accounting for the distinction between the single-minded and the indifferent choice, Sellars (though he was apparently the first to note this distinction) made the mistake of assimilating the intentional modalities to the deontic and was then faced with the puzzle of why there *is* external negation for the latter. Schroeder and Dreier made the converse mistake, in their reading of Gibbard, by effectively assimilating the deontic modalities to the intentional and then puzzling over an apparent *lack* of the two different sorts of negation. It is only when we keep the distinction between deontic and intentional expressions in mind, by distinguishing the single-minded and the indifferent choice for the former and not the latter, that we avoid misconceiving things like the interaction between negation and the various modalities in a plan-theoretic semantics. And that, I have argued, is a matter of seeing the deontic judgment and the collective intention as two species of the genus of practical rationality rather than as species of a genus of one another.

There is another reason for distinguishing the intentional from the deontic mental state. When I introduced the distinction between single-mindedness and indifference in **11.2.2**, I noted some features of these attitudes that suggest the deontic frame of mind involves more cognitive sophistication than the intentional. Among these was the fact that the ability to discriminate an indifferent from a single-minded choice, via the capacity for rejection, requires that one has attitudes toward *other* actions that one *does not* undertake. The intentional state of mind, however, only requires that one be affectively directed toward whatever one decides to do. The ability to discriminate what one rejects from what one rejects rejecting, however, allows one to mark off practical species or determinations from

practical genera or determinables (e.g. attending the colloquium in the front or back of the room and attending the colloquium). Once again it appears that the merely intentional mental state, whether individual or collective, is a cognitively less sophisticated form of thinking than the deontic, and this again supports keeping two sets of books on the plans they involve.

Finally, concerning the question of whether or not the expression ‘not shall’ makes sense in English, or whether it and ‘shall not’ always have the same meaning, I do not think it obvious one way or the other. But the resources of modal logic help clarify the space of options. We can see this by considering analogous cases with the temporal modality ‘will’. Because ‘it will be that A’ and ‘it is not the case that it will be that A’ (using the external negation) are clearly incompatible, we can discern cases where the external negation is not equivalent to an internal negation by looking for contexts where ‘it will be that A and it will not be that A’ (using the internal negation) is coherent. Suppose you are considering coming by my office sometime tomorrow but have not specified when. If you ask me whether I will be in, I might say ‘I will be in my office tomorrow and I won’t be in my office tomorrow’. This sounds sensible where there is an implicit quantification over moments of time, and it does not turn on features of agency. ‘There will be life on Mars and there won’t be life on Mars’ is also acceptable when ‘will’ is understood as quantifying over different moments. Similar uses might be accepted with ‘shall’. Suppose a company is deciding whether to hire a taxi service to take its employees to and from work. When asked whether an employee would use the service, a claim like ‘I shall take the taxi and I shall not take the taxi’, if a bit jarring, seems to make sense in context. As with the temporal modality, its sensibility turns on an implicit quantification over times.

I grant that these examples are not dispositive. But we can make these implicit quantifications explicit, and in temporal logic it is common to treat ‘will’ as a modal operator explicitly quantifying over times either universally (‘will-always’) or existentially (‘will-

sometimes’ – see the discussion at Belnap, Perloff, and Xu pp.32, 152 and 243). Here questions of internal and external negation intersect with the issue of whether a modality has a dual. When explicitly quantifying over temporal moments, ‘will’ operators have duals: e.g., ‘it is not the case that it will-sometimes be that φ ’ and ‘it will-sometimes not be that φ ’ in general have different semantic values, and the strong modal ‘will-always’ can be defined thereby. There appear to be two ways forward, then. If we treat ‘shall’ as either ‘shall-always’ or ‘shall-sometimes’ then ‘shall’ will have both a strong and a weak mode and internally and externally negated shall-sentences will receive different semantic evaluations. On the other hand, we can stick to the simple ‘shall’ locution and let our commitments concerning the metaphysics of agency determine that internal and external negations for shall sentences are equivalent.¹⁰ I choose the latter option, but principally for simplicity. It may be that ‘shall-always’ or ‘shall-sometimes’ are implicitly used in different circumstances, and attention to individual cases might allow one to discriminate features of context, or of the meanings of individual expressions, that serve to specify one or the other uses. For instance, collateral commitments about what we take actions of this sort to involve make it seem that a claim like ‘I shall give to the poor’ either implicitly refers to a particular occurrence of choice (e.g., when I get my next paycheck) or existentially quantifies over circumstances – one can utter a sentence like this without meaning to express commitment to a plan to give to the poor at every point of choice where there is an option to do so. But a sentence like “I shall be more polite to my neighbors” seems to involve universal (or at least general) quantification over circumstances. Features of the meanings associated with the predicates ‘giving to the poor’

¹⁰ It is also possible to eliminate the distinction between internal and external negation while using a quantificational definition for ‘shall’. We do so by requiring that the circumstance parameter for ‘shall’ specifies the moment at which the intention is to be realized, in which case any external negation for a ‘shall-always’ or ‘shall-sometimes’ sentence will be equivalent to an internal negation. Comparison with ‘will’ is again constructive: talk of what is ‘sometimes’ or ‘always’ the case becomes otiose when ‘will’ is evaluated at individual moments, and under a given interpretation two sentences like ‘it is not the case that I will be home at 4:00 pm today’ and ‘I will not be home at 4:00 pm today’ have the same semantic value. See the discussion of the *At-inst* operator in Belnap, Perloff, and Xu (2001), pp.152 and 242-3.

and ‘being more polite to one’s neighbors’ may play a role in determining which sort of reading one gets (this bears comparison with Koons’ discussion of material practical inferences in 2019, chapter 6).

11.3.4 INTENTIONAL HYPERPLANS AND NON-AKRATIC HYPERSTATES

Let an *intentional hyperplan* $h_I \in H_I$ be a maximally determinate plan for what one would choose to do in any circumstance, including what one would do if one were other people. In this case, however, the distinction between indifferent and single-minded choices is irrelevant. That is, an intentional hyperplan h_I is a maximally consistent plan of action such that, for every circumstance C , every agent α able to make a choice at C , and every action A that α is able to choose to perform at C , either (exclusively):

- 1) α chooses to A at C on h_I ; or
- 2) α chooses not to A at C on h_I

The set of *deontic-intentional hyperstates* S_{DI} consists of a set of ordered triples of a world, a deontic hyperplan, and an intentional hyperplan. I use ‘ $\langle w, h_D, h_I \rangle$ ’ to denote elements of S_{DI} . Finally, I ensure that the claim that one shall do something commits one to the claim that one will by semantically evaluating shall-sentences with regard to *non-akratic deontic-intentional hyperstates*, denoted ‘ S_{NADI} ’ where a deontic-intentional hyperstate $s \in S_{DI}$ is non-akratic just in case every choice that is part of the intentional hyperplan h_I of s is such that the action in question will occur in the world of s (I am suppressing a modal interpretation of sentences like *will*(φ), instead treating them as ordinary descriptive sentences):

$s \in S_{NADI}$ iff [$s \in S_{DI}$ and $\forall \alpha \forall A \forall C$ (if α chooses (not) to A in C on h_I of s , then ‘ α will choose (not) to A in C ’ is true at w of s)]¹¹

I use ‘ $\langle w, h_D, h_I \rangle_{NA}$ ’ to denote non-akratic deontic-intentional hyperstates, which I will call ‘deontic-intentional hyperstates’ from here on. One who rejects the implication that to say one *shall* do something is to commit oneself to the claim that one *will* would not use non-akratic hyperstates in the definitions that follow.

11.3.5 A HYPERSTATE SEMANTICS FOR DESCRIPTIVE AND PRESCRIPTIVE SENTENCES, REVISED

The definitions for atomic descriptive sentences and the three kinds of prescriptive sentences can now be given in terms of deontic-intentional hyperstates. For descriptive sentences d :

$[[d]] =_{\text{def.}} \{ \langle w, h_D, h_I \rangle_{NA} : d \text{ is true at } w \}$

‘people are obliged to A in C ’ expresses universally rejecting not doing A in C

$[[\text{people are obliged to } A \text{ in } C]] =_{\text{def.}} \{ \langle w, h_D, h_I \rangle_{NA} : \text{for every } \alpha, \alpha \text{ single-mindedly chooses to } A \text{ in } C \text{ on } h_D \}$

‘people are forbidden to A in C ’ expresses universally rejecting doing A in C

$[[\text{people are forbidden to } A \text{ in } C]] =_{\text{def.}} \{ \langle w, h_D, h_I \rangle_{NA} : \text{for every } \alpha, \alpha \text{ single-mindedly chooses not to } A \text{ in } C \text{ on } h_D \}$

¹¹ Notice that in the antecedent ‘ α chooses (not) to A in C ’ denotes part of a plan and in the consequent ‘ α will choose (not) to A in C ’ denotes a part of a world. Notational variation could disambiguate these two uses, but at the cost of a more barbarous text.

‘people are permitted to A in C’ expresses universally rejecting rejecting doing A in C

$[[\text{people are permitted to A in C}]] =_{\text{def.}} \{ \langle w, h_D, h_I \rangle_{NA} : \text{for every } \alpha, \alpha \text{ either single-mindedly chooses to A in C or indifferently chooses whether or not to A in C on } h_D \}$

11.3.6 INDIVIDUAL AND COLLECTIVE INTENTIONS

As the sentences that give voice to intentions use first-person pronouns, to simplify the presentation I use ‘ ι ’ as a variable taking either ‘I’ or ‘we’ as its value. Against this background, ‘ ι shall A in C’ expresses commitment to the set of deontic-intentional hyperstates where everyone who is one of the group determined by the pronoun chooses to A in C in the corresponding intentional hyperplans:

$[[\iota \text{ shall A in C}]] =_{\text{def.}} \{ \langle w, h_D, h_I \rangle_{NA} : \iota \text{ choose to A in C on } h_I \}$

Similarly, to assert ‘ ι shall not A in C’ is to express commitment to the complement of this set of deontic-intentional hyperstates, which is the set where ι choose not to A in C in the intentional hyperplan of that hyperstate.

Notice that I distinguish the class of agents to which an intention applies from the circumstances in which the intention is realized, rather than building the restriction to the class into the specification of the circumstance as I did with prescriptive sentences. One could accomplish the same result by continuing to use a variable ‘ α ’ that ranges over the unrestricted class of agents and specify that the circumstance in question is one that concerns how either I or one of us would behave, denoted ‘ C_ι ’:

$[[\iota \text{ shall A in C}]] =_{\text{def.}} \{ \langle w, h_D, h_I \rangle_{NA} : \alpha \text{ chooses to A in } C_\iota \text{ on } h_I \}$

As it is useful to keep track of the agents separately from their circumstances of action when dealing with intentions, I shall stick to the former definition.

11.3.7 SOME FEATURES OF A DEONTIC-INTENTIONAL HYPERSTATE SEMANTICS

Various desirable features of the logic of intentions, prescriptions, and descriptions can be established on this semantics. The proofs proceed by showing that the corresponding subset relations obtain among the semantic values of these sentences under an interpretation. For instance, commitment to ‘ ι shall A in C’ commits one to ‘ ι will A in C’:

Proof: $[[\iota \text{ shall A in C}]] \subseteq [[\iota \text{ will A in C}]]$

By the definition of ‘shall’, $[[\iota \text{ shall A in C}]] = \{ \langle w, h_D, h_I \rangle_{NA} : \iota \text{ choose to A in C on } h_I \}$.

Because these triples are non-akratic, it follows that every world in that triple is one where ι will A in C. But that means that the semantic value of ‘ ι shall A in C’ is a subset of the semantic value of ‘ ι will A in C’.

Of course, one may sincerely assert ‘I/we shall A in C’ without oneself or one’s group doing A in C – this is once again a claim about the inferential relations among claims rather than ontological relations among facts. Just so, notice that an assertion of the descriptive claim ‘ α intends to A in C’ (which does not express an intention but reports on its existence), even where ‘ α ’ denotes the speaker, does not commit one to the claim that ‘ α will A in C’. This is reflected in the fact that the semantic value of ‘ α intends to A in C’ is determined by a set of worlds, and these worlds are not in general a subset of the worlds that determine the semantic value of ‘ α will A in C’. At the same time, the claim that one ought to A in C does not commit one to the claim that one will:

Proof: $[[\iota \text{ ought to } A \text{ in } C]] \not\subseteq [[\iota \text{ will } A \text{ in } C]]$

Because deontic sentences are evaluated with regard to different hyperplans than intentional sentences for any given deontic-intentional hyperstate, and because the non-akratic condition is enforced only for intentional hyperplans, the semantic value of ‘ ι ought to A in C ’ will include hyperstates with worlds where ι will not A in C . These worlds are not part of the hyperstates that provide the semantic value for $[[\iota \text{ will } A \text{ in } C]]$, and so $[[\iota \text{ ought to } A \text{ in } C]]$ is not a subset of $[[\iota \text{ will } A \text{ in } C]]$.

Notice that there is also a non-trivial semantic value for ‘ ι shall A in C but ι ought not A in C ’. For $[[\iota \text{ shall } A \text{ in } C]]$ contains *intentional* plans where ι choose to A in C and worlds where ι choose to A in C , while $[[\iota \text{ ought not } A \text{ in } C]]$ contains *deontic* plans where ι single-mindedly choose not to A in C . Because the semantics of the former imposes no restriction on the deontic plans determined by the latter, and vice versa, there will be deontic-intentional hyperstates at the intersection of these two sentences. It is true that in such a mental state one is in conflict with oneself over what one *intends* to do and what one thinks one *should* do, as exemplified by the fact that there is no decision one can make that satisfies both sets of plans. But alas, that is a feature of the mental life of sublunary beings. Thus akrasia is possible even though the hyperstates that model akratic thoughts are non-akratic in the intention position. Furthermore, concerning collective intentions, if we were to follow Sellars and interpret ‘ought’ in terms of ‘shall_{we}’ it would be hard to make sense of someone who said both ‘we ought to A ’ and ‘we shall not A ’. But there are clearly cases where one can hold that a collective of which one is a member is about to do something which one thinks they ought not.¹² At the same time the semantic value of a complex strong deontic, weak deontic, and

¹² Reflection on this point first led me to distinguish the deontic from the intentional hyperplan, though I have since become convinced that doing so, and the distinction between single-mindedness and indifference in the former, is central to an adequate understanding of practical rationality.

intentional sentence like ‘I ought to go to the lecture, I may sit in the front, I may sit in the back, and I shall sit in the front’ is well-defined as the intersection of the semantic values of the separate sentences:

$$\begin{aligned} & [[\text{I ought to go to the lecture, I may sit in the front, I may sit in the back, and I shall sit in} \\ & \text{the front}]] = [[\text{I ought to go to the lecture}]] \cap [[\text{I may sit in the front}]] \cap [[\text{I may sit in the} \\ & \text{back}]] \cap [[\text{I shall sit in the front}]] = \{ \langle w, h_D, h_I \rangle_{NA}: \text{I single-mindedly choose to go to the} \\ & \text{lecture on } h_D \} \cap \{ \langle w, h_D, h_I \rangle_{NA}: \text{I either single-mindedly choose to sit in the front or} \\ & \text{indifferently choose whether or not to sit in the front on } h_D \} \cap \{ \langle w, h_D, h_I \rangle_{NA}: \text{I either} \\ & \text{single-mindedly choose to sit in the back or indifferently choose whether or not to sit in the} \\ & \text{back on } h_D \} \cap \{ \langle w, h_D, h_I \rangle_{NA}: \text{I choose to sit in the front on } h_I \} \end{aligned}$$

That this is a practically rational state of mind is reflected in the fact that, unlike the state of mind expressed by ‘I ought to A in C but I shall not’, there is a choice one can make that satisfies all of these sets of plans. For the sets at the intersection of the middle two conjuncts rule out those deontic-intentional hyperstates where I choose, on the deontic hyperplan h_D of the corresponding hyperstates, either single-mindedly to sit in the front or single-mindedly to sit in the back, and this leaves only those hyperstates whose deontic hyperplans have me choosing indifferently to sit in either the front or the back. In such a decided state the only way to conform both to one’s deontic and to one’s intentional plans is to choose indifferently to sit in the front – to do so, then, is to live up to the commitments one has expressed in making this complex deontic/intentional claim.

Finally, the GENUINE AGENCY condition, and the fact that intentional hyperplans do not mark indifference and single-mindedness, ensures that the semantic values for the internal and the external negations of intentions are equivalent:

Proof: $[[\text{it is not the case that } \iota \text{ shall A in C}]] = [[\iota \text{ shall not A in C}]]$

By the definition of ‘shall’, $[[\iota \text{ shall A in C}]] = \{ \langle w, h_D, h_I \rangle_{NA} : \iota \text{ choose to A in C on } h_I \}$.

Call this set X. The semantic value of the negation of ‘ ι shall A in C’ is the complement of X. It is a feature of agency, encoded in GENUINE AGENCY, that one can only choose to do something at a circumstance where one could also choose not to do it—otherwise the choice is no choice at all. By the definition of hyperplans, then, if an agent does not choose to A in some circumstance C where she could, then the agent chooses *not* to A in C. Because intentional hyperplans concern only what the agent chooses to do and not do, it follows that the complement of X is $\{ \langle w, h_D, h_I \rangle_{NA} : \iota \text{ choose not to A in C on } h_I \}$. But this latter set is the semantic value of ‘ ι shall not A in C’.

By interpreting ‘shall’ in terms of either ‘shall-always’ or ‘shall-sometimes’ this equivalence would not hold.

11.3.8 A NOTE ON BRATMAN’S THEORY OF COLLECTIVE INTENTIONS

A more extended comparison of this theory of practical rationality with current work in the literature on collective intentionality is deferred for another time. But there is one aspect of Michael Bratman’s proposal that this account nicely accommodates, and it does so in a way that illuminates a central commitment of Sellars about the identity and individuation of collective intentions. In various places Bratman discusses the norms of rationality he associates with the planning states that lie at the back of collective intentions (cf. 2006, 2014, the Introduction to 2018, and the concise statement of his position in 2017), and he contends that an account of collective intentionality should also satisfy a condition of rational adjustment on these planning states as they develop over time. From (2006, p.3):

[M]y intention that we *J*...imposes rational pressure on me, as time goes by, to fill in my subplans in ways that fit with and support yours as you fill in your subplans; and vice versa.

In defending the contention that rational pressure forces individuals who share a collective intention to fill in their individual plans in such a way as to realize that collective intention, Bratman writes (2010, p.8):

A fundamental ground of these human capacities for temporally extended and shared intentional agency are human capacities for planning agency. In saying that these planning capacities are a fundamental ground, I mean that the proper exercise of these planning capacities, given relevant contents of the plans, relevant contexts, and relevant inter-relations with past, future, and others, will realize phenomena of temporally extended and/or shared intentional activity.

On the hyperstate analysis provided here, the Boolean logic of set theory goes some way toward explaining how the norms of rationality associated with planning mental states enable those mental states to “realize phenomena of temporally extended and/or shared intentional activity”.

This can be seen with an example. Let $S_\alpha \subseteq S_{NADI}$ denote the set of deontic-intentional hyperstates that models the content of the mental state of the agent α . Suppose Alice and Bertolt share a collective intention to get a beer after the conference, but that no preferences for how to get to the pub have been specified (I am suppressing considerations of cognitivism – the question of whether those who share a collective intention must believe or know that

they do so). Setting ‘ α ’ as ‘Alice’, ‘ β ’ as ‘Bertolt’ and ‘C’ as ‘after the conference’ this mental state is modelled by the hyperstate S_1 such that:

$$S_1 = (\{ \langle w, h_D, h_I \rangle_{NA} : \alpha \text{ chooses to get a beer with } \beta \text{ in } C \text{ on } h_I \} \cap \{ \langle w, h_D, h_I \rangle_{NA} : \beta \text{ chooses to get a beer with } \alpha \text{ in } C \text{ on } h_I \})$$

Now, let us suppose that Bertolt proposes to Alice that they walk to the pub together, and that in doing so he comes to hold his part of the collective intention that they do so. Bertolt’s mental state (denoted ‘ S_β ’) is now more restricted than S_1 , imposing as it does the additional requirement that Alice and Bertold walk together to the pub:

$$S_\beta = S_1 \cap (\{ \langle w, h_D, h_I \rangle_{NA} : \alpha \text{ chooses to walk with } \beta \text{ to the pub in } C \text{ on } h_I \} \cap \{ \langle w, h_D, h_I \rangle_{NA} : \beta \text{ chooses to walk with } \alpha \text{ to the pub in } C \text{ on } h_I \})$$

For Alice and Bertold to share this collective intention, Alice will likewise need to restrict her hyperstates to those that include the collective intention to walk to the pub. On this account, the rational pressure to fill in one’s individual plans to satisfy a collective intention (and the explanation of this phenomenon of collective intentionality) is reflected in the fact that more precise collective planning states are given by the intersection of the original states with the planning states that arise as a collective plan becomes more determinate. Finally, notice that this semantics for collective intentions vindicates a claim made by Sellars (1968, p.217):

“The *intendings* are two in number, but the *content* of these intendings is the *same*”. Where Alice and Bertold share a collective intention, and where collective intentions are individuated by the individuals possessing them whereas their content is individuated by hyperstates, Alice and Bertold will have numerically distinct collective intentions with the same content.

11.4 SUMMING UP AND LOOKING AHEAD

In this essay I have provided a hyperstate semantics sufficient to distinguish the descriptive, prescriptive, and intentional contents of various sentences and the mental states their assertions give voice to, and I have examined and put to criticism some of Sellars' views in this vicinity. In doing so I hope to have underwritten my contention that Sellars deserves to be placed alongside the 'big four' of collective intentionality as one of the significant figures in this tradition, and that Sellars' views are deserving of more consideration in the literature on collective intentionality. Beyond its historical relevance, this semantics offers the possibility of a common language for further investigation into issues of rationality, intentionality, and our shared social spaces. For on the view developed here practical reasoning, moral reasoning and individual and collective intentions can all be understood through a common framework that handily accommodates various features of theoretical and practical rationality. It also, I have suggested, sheds light on the nature of practical rationality as it is manifest in deontic and intentional cognition.

APPENDIX: *DE RE* AND *DE DICTO* MODALITY

Consider the *de dicto* claim 'it ought to be that every K chooses to A in C' (e.g., it ought to be that anyone who is a member of my department attend the lecture). We formalize this claim with kind-restricted quantification as follows:

$$\text{Ought}(\forall K, \alpha)(\alpha \text{ chooses to } A \text{ in } C)$$

Intuitively, this says that, no matter who the Ks actually are, anything that *were* a K ought to A in C. The *de re* claim 'every K ought to choose to A in C' (everyone who is a member of

my department ought to attend the lecture), by contrast, says that every *actual* K is such that it ought to A in C:

$$(\forall K, \alpha) \text{Ought}(\alpha \text{ chooses to A in C})$$

To discriminate the semantic values of *de re* and *de dicto* occurrences of the deontic modals, it is correspondingly necessary to relativize semantic values to points of evaluation. Each point of evaluation assigns an object as an extension to every term, a set of objects as the extension of every quantifier, an n-tuple of objects to every n-place predicate, etc. In order to evaluate *de re* modal claims, principles of identity and individuation are then needed in order to track individuals across different points of evaluation. Consideration of possible-worlds analyses helps illustrate the point. In a possible-worlds semantics the *de dicto* claim is true at a world w just in case every morally-best world w_1 accessible from w is such that the following holds: every K at w_1 chooses to A in C at w_1 . In this case one only looks at how otherworldly Ks behave at those other worlds. By contrast, the *de re* claim is true at a world w just in case every morally-best world w_1 that is accessible from w is such that the following holds: every K at w chooses to A in C at w_1 . *De re* modality thus involves a principle of transworld identification – to evaluate *de re* occurrences of a modal we have to consider what the actual Ks (the individual Ks at the point of evaluation) are doing at other worlds.

Even though a hyperstate semantics makes use of plans to evaluate the contents of prescriptive and intentional mental states, in addition to worlds, the difference between *de re* and *de dicto* deontic modality can be modelled on a hyperstate semantics if we relativize the semantic evaluation of sentences to worlds alone. Letting $[[\varphi]]_w$ denote the semantic value of φ at a world w , and where '@' denotes the world of evaluation (in the usual case, the actual world), the *de dicto* occurrence of the strong deontic modality is interpreted as follows:

$[[ought(\forall K, \alpha)(\alpha \text{ chooses to } A \text{ in } C)]]_{@} =_{\text{def.}} \{ \langle w, h_D, h_I \rangle_{NA} : \text{every } K \text{ in } w \text{ single-mindedly chooses to } A \text{ in } C \text{ on } h_D \}$

The *de re* occurrence receives the following interpretation:

$[[(\forall K, \alpha)ought(\alpha \text{ chooses to } A \text{ in } C)]]_{@} =_{\text{def.}} \{ \langle w, h_D, h_I \rangle_{NA} : \text{every } K \text{ in } @ \text{ single-mindedly chooses to } A \text{ in } C \text{ on } h_D \}$

De re and *de dicto* talk of permissions and prohibitions can be understood in the same manner.

Notice that just as the interpretation of *de re* modality on a possible-worlds semantics requires tracking individuals from the actual world to various possible worlds, so does *de re* modality require tracking agents from the actual world to the plan-elements of various hyperstates. But on a hyperstate semantics a principle of identifying and individuating individuals across worlds and plans is already needed for *de dicto* modality, and in future work I hope to examine such principles in more detail.

Acknowledgements

Work on this chapter was supported by the joint Lead-Agency research grant between the Austrian Science Foundation (FWF) and the Czech Science Foundation (GAČR), *Inferentialism and Collective Intentionality*, GF17-33808L. Portions of these ideas have been presented at various conferences and workshops over the last two years, and I owe particular thanks for discussion with Bill deVries, Ulf Hlobil, Jeremy Koons, Ladislav Koreň, Cathy Legg, Ronald Loeffler, Jim O'Shea, Jarda Peregrin, Mark Risjord, Carl Sachs, and Luz Christopher Seiberth.

BIBLIOGRAPHY

- Belnap, N., M. Perloff, and M. Xu (2001). *Facing the Future: Agents and Choices in Our Indeterminist World*. New York: Oxford University Press.
- Brandom, R. B. (1994). *Making It Explicit: Reasoning, Representing, and Discursive Commitment*. Cambridge: Harvard University Press.
- Bratman, M. (1993). Shared Intention. *Ethics* 104(1), 97-113.
- (2006). Dynamics of Sociality. *Midwest Studies in Philosophy* XXX, 1-15.
- (2010). Agency, Time, and Sociality. *Proceedings and Addresses of the American Philosophical Association*, 84(2), 7-26.
- (2014). *Shared Agency: A Planning Theory of Acting Together*. New York: Oxford University Press.
- (2017). Rational Planning Agency. *Royal Institute of Philosophy Supplement*, 80, 25-48.
- (2018). *Planning, Time, and Self-Governance: Essays in Practical Rationality*. New York: Oxford University Press.
- Chant, S.R., Hindriks, F., Preyer, G. (Eds.). *From Individual to Collective Intentionality: New Essays*. New York: Oxford University Press.
- Dreier, J. (2006). Negation for Expressivists: A Collection of Problems with a Suggestion for their Solution. In Russ Shafer-Landau (Ed.), *Oxford Studies in Metaethics, Volume 1* (pp.217-233). New York: Oxford University Press.
- (2009). Relativism (and Expressivism) and the Problem of Disagreement. *Philosophical Perspectives*, 23, 79-110.
- Gibbard, A. (2003). *Thinking How to Live*. Cambridge: Harvard University Press.
- Klemick, G. (2018). Sellars' Metaethical Quasi-Realism. *Synthese*, published online 18 May, 2018.
- Koons, J. (2019). *The Ethics of Wilfrid Sellars*. New York: Routledge.
- (**This Volume**).
- Ludwig, K. (2016). *From Individual to Plural Agency: Collective Action, Volume 1*. New York: Oxford University Press.
- Mele, Alfred R. (1989). "Intention, Belief, and Intentional Action" *American Philosophical Quarterly* 26 (1): 19-30.
- Olen, P., Turner, S. (2015). Durkheim, Sellars, and the Origins of Collective Intentionality. *British Journal for the History of Philosophy*, 23 (5), 954-975.
- Peregrin, J. (This Volume)**.
- Schweikard, D.P, Schmid, H.B. (2013). Collective Intentionality. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*. <https://plato.stanford.edu/entries/collective-intentionality/>.
- Schroeder, M. (2008). *Being For: Evaluating the Semantic Program of Expressivism*. New York: Clarendon Press.
- Silk, A. (2015). How to Be an Ethical Expressivist. *Philosophy and Phenomenological Research* XCI (1), 47-81.
- Sellars, W. (1951). Obligation and Motivation. *Philosophical Studies*, 2(2), 21-25.
- (1953). Inference and Meaning. *Mind*, 62 (247), 313-338.

- (1963). Imperatives, Intentions, and the Logic of “Ought”. In H-N. Casteñeda and G. Nakhnikian (Eds.), *Morality and the Language of Conduct* (pp. 159-214). Detroit: Wayne State University Press.
 - (1966a). ‘Ought’ and Moral Principles. Transcribed into hypertext by Andrew Chrucky in June 2005 from a copy saved by Willem deVries.
<http://www.ditext.com/sellars/omp.html>. Accessed May 9, 2015.
 - (1966b). Thought and Action. In K. Lehrer (Ed), *Freedom and Determinism* (pp.105-139). New York: Random House.
 - (1967a). Form and Content in Ethical Theory (the 1967 Lindley Lecture). Department of Philosophy, University of Kansas.
<https://kuscholarworks.ku.edu/bitstream/handle/1808/12383/Form%20and%20Content%20in%20Ethical%20Theory-1967.pdf>. Accessed May 9, 2018.
 - (1967b). “Science and Ethics”. In Wilfrid Sellars, *Philosophical Perspectives: Metaphysics and Epistemology*. Atascadero: Ridgeview Publishing Company. Pp. 191-209.
 - (1968). *Science and Metaphysics: Variations on Kantian Themes*. Atascadero: Ridgeview Publishing Company (1992).
 - (1969). Are there Non-Deductive Logics? In N. Rescher (Ed.), *Essays in Honor of Carl G. Hempel* (pp.83-103). Dordrecht: D. Reidel Publishing Company.
 - (1976). Volitions Re-Affirmed. In M. Brand and D. Walton (Eds.), *Action Theory* (pp.47-66). Dordrecht: D. Reidel Publishing Company.
 - (1980). On Reasoning about Values. *American Philosophical Quarterly*, 17(2), 81-101.
- Stovall, P. (Forthcoming). Normative Attitudes and Collective Intentions: The Ground of Brandomian Inferentialism. *The Social Institution of Discursive Norms*. New York: Routledge.
- (Manuscript). A Unified Semantics for a Language of Theoretical and Practical Reasoning.
- Tuomela, R. and Miller, K. (1988). We-Intentions. *Philosophical Studies*, 53, 367-389.