THE METAPHYSICS OF PRACTICAL RATIONALITY
INTENTIONAL AND DEONTIC COGNITION[1]

Preston Stovall
University of Hradec Králové

Despite growing appreciation in recent decades of the importance of shared intentional mental states as a foundation for everything from divergences in primate evolution, to the institution of communal norms, to trends in the development of modernity as a socio-political phenomenon, we lack an adequate understanding of the relationship between individual and shared intentionality. At the same time, it is widely appreciated that deontic reasoning concerning what ought, may, and ought not be done is, like reasoning about our intentions, an exercise of practical rationality. Taking advantage of this fact, I use a plan-theoretic semantics for the deontic modalities as a basis for understanding individual and shared intentions. This results in a view that accords well with what we currently have reason to believe about the phylogenetic and ontogenetic development of norm psychology and shared intentionality in human beings, and where original intentionality can be understood in terms of the shared intentionality of a community.

Keywords: collective intentionality, deontic modality, practical rationality

I Shared Intentionality as a Theme of Modern Philosophy

One way of telling the story of modernity, as a social or political phenomenon, is that it involved a dawning recognition of the power of shared or collective intentional agency as a means of shaping who we are and how we see the world. As a tour through modern philosophy, this telling might include discussion of Rousseau's conception of the General Will (cf. Miller and Tuomela 2014, and Tuomela 2013: 94-6), Kant's understanding of the Categorical Imperative as the demand to constitute a universal kingdom of ends (1996: 83ff.), Hegel's claim that "the experience of what Spirit is" arises as the recognition of an "'I' that is 'we' and a 'we' that is 'I'" (1977: 110), Peirce's view that the identities of individual persons

have their ground in the universal community of inquiry (cf. Burgh et al. 2006 chapter 2), and Wilfrid Sellars' effort, which he understood as working out ideas in Kant and Peirce, to characterize the moral frame of mind as a species of shared intentionality (e.g. 1951, 1963, 1968, 1976). Prominent contemporary examinations of the ground of social reality in shared intentionality include Gilbert (2014), Searle (1995, 2010), Tuomela (2007, 2013), and the essays in Schmitt (2003a), Schmitz (2013), and part VI of Jankovic and Ludwig (2018). Owing to increased interest over shared intentionality in recent Anglophone philosophy, and in related areas of evolutionary anthropology, developmental psychology, and the social sciences, a story of this sort offers the prospect of relating major trends in the modern period to contemporary debates about human thought and agency in the sciences and the humanities, by using developments in one field to shed light on work in others.

In the contemporary philosophical literature on shared intentionality, one of the recurring questions concerns whether and in what sense shared intentions can be reduced to individual intentions (for representative discussions see Birch 2018, Bratman 2014 chapters 1 and 6, Ludwig 2016 chapters 9-12, Schweikard and Schmid 2013 section 3, and Tuomela 2013 chapter 2). The effort to provide some such reduction is often connected to a felt need to avoid reifying collective minds as bearers of mental states. More-or-less overt ontological reification occurs across much of this literature; the full passage from Hegel partially quoted above, for instance, reads:

What still lies ahead for consciousness is the experience of what Spirit is—this absolute substance which is the unity of the independent self-consciousnesses which, in their opposition, enjoy perfect freedom and independence: 'I' that is 'we' and 'we' that is 'I'.[2]

---

[2] Was für das Bewußtsein weiter wird, ist die Erfahrung, was der Geist ist, diese absolute Substanz, welche in der vollkommenen Freiheit und Selbstständigkeit ihres Gegensatzes, nämlich verschiedener für sich seiender Selbstbewußtsein, die Einheit derselben ist; *Ich*, das *Wir*, und *Wir*, das *Ich* ist.

To disarm the apparent reference to collective minds that talk of shared or collective mental states seems to involve, philosophers often try to reduce any apparent mindedness in a collective to the mindedness of its members, trading the tendency toward what Schmitt (2003b) calls *supraindividualism* for a tendency toward *individualism*. Meanwhile, some philosophers hold that even if there are no irreducibly collective *minds* there are irreducibly collective *mental states* (cf. the discussion of we-mode intentionality in Tuomela 2007 and 2013, and Miller 2014). The bearers of these mental states may be individuals, but as mental states they are essentially collective (views that tolerate more supraindividualist perspectives include Gilbert 1989 and 2000, List and Pettit 2011, Pettit 2003, Schmid 2018, Schmitt 2003b, and Searle 1995).

Despite the variety in their views, two commitments common across much of this literature contribute to worries about supraindividualism. First, philosophers tend to model theories of shared intentionality on theories of individual intentionality: starting with an antecedent understanding of the latter, the capacities thought operative there are repurposed to account for shared intentions (e.g. Bratman 2014: 14 and 128-30, and Ludwig 2018: 477). Chant (2018) calls this the "wash, rinse, and repeat" (WRR) approach. From p.14:

> Employing WRR as a method for analyzing group-level concepts such as collective intention is to imply a kind of isomorphism between theories of individual action and theories of collective action. The isomorphism is between the respective sets of individual- and group-level concepts (e.g. "individual agent" and "collective agent"), and the explanatory relationships between them. Taking on WRR is tantamount to assuming that if there is an explanatory relationship between two individual-level concepts, there should be a corresponding pair of collective-level concepts bearing the same explanatory relationship to each other.

Alongside this first framing commitment, debates about shared minds have been shaped by representational analyses provided for sentences like "we built the house together" and "the homeowner's association is responsible for the change in policy". These sentences are thought to be true just in case certain conditions hold, and they are semantically evaluated as representations of these truth conditions (e.g. in Bratman 2014 chapter 4, Ludwig 2007, 2014, and chapters 12 and 13 of 2016, and Tuomela and Miller 1988). We gain a rather different view of shared intentionality if we trade these two commitments for, in the first instance, an *analogical* rather than *isomorphic* understanding of shared intentionality, and in the second instance an analysis of the sentences that give *expression* to shared intentions rather than those that purport to *report on* or *represent* their existence.

The first shift might not seem that significant. Although Chant characterizes WRR as an isomorphic understanding of shared intentions on the basis of an understanding of individual intentions, it is not uncommon to see philosophers frame their accounts in terms of analogical relationships between the two domains (e.g. Bratman 2014: 10 and 30-1, and Tuomela 2007: 140). For this reason, I think we should understand Chant as diagnosing a tendency or trend in recent philosophy rather than an explicit or universally exemplified commitment. The important point, for both her and myself, is that some of the troublesome ontological puzzles philosophers have faced are artifacts of the modelling methods prevalent in the contemporary literature. I will argue that we can solve, or dissolve, some of these puzzles by thinking more explicitly about the analogical character of our models.

Work in the philosophy of science over the last half-century has equipped us with a set of protocols for the critical application of analogical reasoning in theory construction (cf. Bartha 2010, Genter 1983, Gentner, Holyoak, and Nakinov 2001, and Hesse 1952 and 1966). In its general form, we use an antecedent understanding of one domain of inquiry (the analogical base) for framing an understanding of a different domain (the analogical target).

After identifying hallmark features of the analogical base and target that are situated in similar networks of operation (causal, geometrical, conceptual, etc.), we give a commentary on the sorts of relations in the base that can be inferred to hold among the correlate relations in the target. While there must be some structural similarity between the analogical base and target if the analogy is to get off the ground—consider the sun at the center of the solar system and the nucleus at the center of the atom—much of the explanatory value of analogical reasoning derives from the commentary one gives concerning which inferences can be drawn from base to target. The solar system model of the atom concerns geometrical relationships, and in the commentary we block the inference that nuclei irradiate electromagnetic energy on the basis of gravity-induced nuclear fusion. Just so, an analogical understanding of shared intentionality may include a commentary sufficient to defuse it of the charge that it reifies the group into some kind of supraindividual mind. Metaphysical concerns may be avoided from the beginning, by simply not importing the relevant explanatory resources from base to target.

Analogical reasoning has the additional benefit of helping reframe our understanding of a subject matter that may be otherwise difficult to comprehend. Using an initial model drawn from the base, we investigate the analogical target to see whether the hypothesized relations posited from our reading of the base are present there. Over the course of the investigation, we update the model according to what we discover about the target domain. As a consequence of this revisionary process, an analogical conception of the relationship between two domains may reorient how we understand the very things whose metaphysical status we dispute. In one of its forms, analogical reasoning about two different domains allows us to construct a more general category that subsumes them both (cf. general theories of teleological systems derived from Darwinian accounts of biology; for discussion see Stovall 2015). This capacity to change our understanding of a target of inquiry led Hegel to associate inference by analogy with the syllogistic figure Peirce would label 'hypothesis' and

subsequently assimilate to abduction, and both philosophers thought this form of reasoning was at the core of scientific and philosophical inquiry—for Hegel, it involved the ability to reason with universal terms as middle terms in a syllogistic inference, and so it constituted our grasp of universality as the most general kind of thinking (see Redding 2003 and Stovall Forthcoming A). They also thought this form of reasoning was going to be instrumental for the advance of intellectual and social projects rooted in the modern period, and at the end of the essay I will suggest that we heirs of modernity can still profit from this way of thinking.

Concerning the second shift in commitment I take on here, recent debates have been unduly influenced by the thought that the mind's core task is to represent the world. Much of what we do is not strictly representational but agency-involving, and meant to change the world instead of describe it. Representational capacities are involved in action, of course, but that is not so say a claim like "we are going to the park this afternoon" or "you should be more polite to your neighbors" purports to represent a fact—it may instead be to rehearse an intention, or a call for another to do the same. This distinction has its correlate in the organism. Sensory and motor neural pathways are physiologically distinct structures, each carrying its own intentionality and direction of fit: cognitive operations of the former are successful or correct when the mind 'fits' or corresponds to the world, whereas the operations of the latter are successful when the world fits the mind. Cases of mind-world fitness may be conceived in exclusively representational terms with a tolerable degree of accuracy, but an adequate understanding of world-mind fitness requires the categories of practical rationality. And we can recognize this point without losing sight of the fact that we in some sense represent intentional facts when we say, e.g., that we built the house together.

At any rate, the philosopher is not required to begin studying shared intentionality through a truth-conditional analysis of sentences that appear to report on the existence of shared intentions. Problems of the apparent existence of superminds are more pronounced if

we orient ourselves along the representational and truth-conditional dimensions of linguistic activity. Instead, if we are interested in shared intentional states of mind we should pay attention to the sentences used to give voice to shared intentions. Consequently, rather than propose truth conditions for sentences like 'we built the house together' I will offer an analysis of the semantic content of sentences like 'we *shall* build the house together', where the modal term 'shall' connotes the expression of an intention rather than a claim about its existence. This is a shift from a truth-conditional mind-world analysis of shared intentions to a plan-conditional world-mind analysis. More specifically, I propose we understand shared intentional mental states by analogy with the mental states associated with deontic judgments giving voice to what people ought, may, and ought not do, and I will use a planning semantics to understand both the intentional and the deontic modalities. This analogy is motivated by the observation that practical rationality is operative as a hallmark feature of both kinds of cognition, and my claim is that these frames of mind alike involve planning on how one would behave *were one any member of the relevant community*.

II Background and Outline[3]

I want to begin by addressing two sets of worries one might have at the outset. First, it might be thought that the use of logical tools as a basis for investigating the nature of cognition was the resuscitation of a kind of metaphysical speculation that had long died out, to the benefit of our intellectual and scientific ecosystem. Here I want to emphasize two things and encourage the reader to keep them in mind in what follows. First, the model presented below is probative and subject to revision. I do not supply the final answers on what I discuss, and I try to use the formalism as a tool to help clarify what seem like otherwise plausible views about human cognition. Second, the empirical study of human cognition is at least as relevant

---

[3] My thanks to a reviewer for suggesting this discussion.

a source of model construction as the formal details that motivate my commitments here, and I strive to keep the discussion consonant with and informed by empirical work (more on this in a moment). As a second concern, one might worry that in an essay of this length the treatments of these ideas will be too quick and superficial to carry the full weight of the claims placed on them. Here I concede guilt while appealing to material that develops the view elsewhere.[4] And because this essay is a more-or-less philosophical investigation into the relationship between intentionality and deontic cognition, the social-scientific and evolutionary background is kept out of view entirely. In part to allay worries that my approach here is either too formalist or too superficial, then, it may do to say something about that background now.

In the last three decades, the empirical study of shared intentionality has developed into a burgeoning field of research across the behavioral, psychological, social, and evolutionary sciences. This work has begun to frame a picture of ourselves as paradigmatically norm-tracking, norm-enforcing, and collectively intentional animals (see Schmidt and Rakoczy 2018, and Tomasello, et al. 2005 for overviews). While there is some evidence for coordinated activity that might be thought of as shared intentionality in non-human animals, from an early age human beings exhibit a range of shared and norm-enforcing behaviors unlike what is generally seen in even our closest primate relatives: before their first birthday human infants appear to understand human-like movements as goal directed, and around their first birthday they begin to share attention and emotions with other humans (Tomasello 2005: 688-9); between 12 and 18 months human infants show a tendency to

---

[4] This essay results from an interdisciplinary project, underway for the last three years at two universities, investigating the shared-intentional ground of rule-governed discursive cognition as a part of the natural world. In other published work I critically evaluate Sellars' assimilation of deontic cognition to shared intentionality (Forthcoming B), and I argue that the ability to speak a natural language presupposes a developmental progression from the intentional to the deontic frame of mind (Forthcoming C). In a manuscript (Forthcoming D) I develop these ideas more extensively, relating them both to existing views in the philosophical literature and to ongoing empirical work on shared intentionality in evolutionary anthropology, neurophysiology, and the social sciences.

engage in both instrumental and playful cooperation (Schmidt and Rakoczy 2018: 698); by the age of 3 children distinguish a range of features of the normative landscape, including the distinction between conventions and moral norms (where the former depend on the agreement of the community; see Chudek and Henrich 2011); by this age they also recognize that the entitlement to do something induces commitments of non-interference by others (Schmidt et al. 2013), and they enforce norms on others (Rakoczy et al. 2008); between the ages of 3 and 5 children gain the ability to better cooperate in a shared activity by swapping roles with other people, an ability that chimpanzees appear not to possess (Fletcher et al. 2012); and by the age of 5 they spontaneously create and enforce norms when interacting with others (Göckeritz et al. 2014). So it appears human cognition is intimately connected to both deontic judgment and shared intentionality. From a physiological standpoint, there is good reason to treat these forms of cognition as involving multiperspective planning states. In the 1990s the motor-representation complex of the central nervous system in macaque monkeys was shown to be operative *both* when the monkey performs various activities *and* when the monkey observes others perform them (see Rizzolatti and Sinigaglia 2010 for an overview). It is as if the observer monkey was putting itself into the position of the observed and preparing to act from that point of view. Since then, we have discovered that this system is present across much of the animal kingdom. Work on the mirror system went through a period of intense popularization in the philosophy of mind and the philosophy of science during the first decade of the third millennium, and it has occasionally been put to use in philosophical accounts of shared intentionality as well (Butterfill 2015 and 2018, and Tuomela 2007: 69-74, 268 fn.35, and 297 fn.26).

The formalism I defend is informed by the view that norm psychology and shared intentionality are characteristic features of human cognition, and that shared intentionality is partially realized, in the wetware of the brain, in patterns of multiperspective neural

coordination across a community. I say 'partially' because mirroring in the motor-representation complex is at best a (naturally) necessary rather than a sufficient condition for some kinds of shared intentionality: it was perhaps a condition on its evolution in human beings, not something that must be present in every case of shared intentionality. For many animals exhibit motor-representational mirroring, but shared intentionality appears to be very rare in the animal kingdom—so mirroring is not sufficient for shared intentionality. And it is not necessary, as we can settle a business decision by conference call or decide as a nation to go to war with very little in common across our motor-representational complexes. Instead, it seems as though the coordinated neural mirroring that characterizes much of the animal kingdom interacts, in human ontogeny, with a suite of more cognitively complicated capacities peculiar to human beings, resulting in our capacity to adopt the shared affective, cognitive, and practical stances that characterize so much of human life (cf. Butterfill 2018). I suspect that shared intentionality in human beings is best conceived as the product of a process of evolutionary and educational development that includes not only neural mirroring but also the human proclivity to share attention, emotion, and a range of other affective/motivational states.

While I leave this discussion behind in what follows, the central hypothesis suggested by the analogical reasoning below is meant to contribute to the empirical study of shared intentionality and norm psychology in human beings. It does so by proposing one sense in which the form of cognition associated with the latter is more sophisticated than that of the former. Toward that end, my terminology and framework is Michael Tomasello's (2014). Tomasello argues for a two-stage process of evolution from early hominids to modern humans: a stage of *joint* intentionality where dyads and small groups could share intentionality over spatio-temporally proximate kinetic behavior, and a stage of *collective* intentionality where an expanded capacity to share intentionality allowed for the sense of

oneself as the member of an abstract community. At this second stage, Tomasello contends, we have the (in-principle) universal sense of community that conditions our understanding of the *true* and the *moral* as that which constrains the thought and action of everyone. Tomasello notes that there may be additional stages between these two transition points from brute to civilization (2014: 151), and I can be read as positioning an intermediate stage in Tomasello's framework, between the development of collective intentionality and the capacity to think about what is true and moral. I do so by using the analogical reasoning below as a basis for re-examining some of the presuppositions of recent philosophical work on shared intentionality, and I defend the following hypothesis: the intentional state of mind is a more physiologically primitive form of practical rationality than that which is characteristic of the deontic state of mind, so that intentional mental states are prior in the order of being to deontic mental states. More specifically, I show that the need to distinguish the strong and weak deontic modal forces (e.g. *ought* and *may*) requires appealing to the notion of a *single-minded* choice attitude, which involves bearing attitudes *both* toward what one chooses *and* toward choices that are incompatible with what one chooses. An animal able to think single-mindedly could then adopt an *indifferent* frame of mind as well, according to the notion of indifference defined below. And with the choice attitudes of single-mindedness and indifference it is possible to model the planning mental states associated with the strong and weak deontic modalities. The intentional modalities, by contrast, do not distinguish strong and weak modal forces, and so we can model the sentences giving voice to individual and shared intentions with planning states that make no appeal to the capacity to bear attitudes both toward what one does and toward choices incompatible with what one does. The conclusion I draw—adding detail to Tomasello's philosophical anthropology, and presented as a hypothesis for further examination—is that the transition from the joint intentionality of small groups to the truth- and morality-mongering intentionality of modern humans required

the development of a more sophisticated kind of practical cognition (the single-minded thinking that comes with deontic reasoning) than is necessary for the exercise of shared intentionality, whether joint or collective. This fits remarkably well with a claim made in the closing paragraph of the review article by Schmidt and Rakoczy (2018: 698, emphasis added):

> So, one picture that is worth being explored more systematically in future research is that while humans and other species, notably primates, share basic forms of individual intentionality (and the corresponding natural norms of correctness and success), uniquely human forms of norm psychology and uniquely human forms of shared intentionality develop in close tandem in early ontogeny, *the former building on and growing out of the latter.*

At the end of the essay I offer a second hypothesis suggested by this model: we can understand individual intentions on the basis of our understanding of shared intentions rather than vice versa, so that while shared intentional mental states might be prior in the order of *being* to deontic mental states, shared intentional mental states are prior in the order of *understanding* to individually intentional ones. In a sense to be explicated in below, original intentionality is the shared intentionality of a community. That this is a theme resonant with certain ideas in modern philosophy will remain mostly undiscussed.

III Understanding the Intentional Modalities by Analogy with the Deontic Modalities
Though not equivalent in all contexts, I will indifferently speak of sentences under an interpretation, of utterances, assertions, judgments, or claims, of mental states associated with utterances, and of propositional commitments.

It is widely accepted that shared intentions, as a species of intention, are or involve plans of action (Bratman, e.g. 2014, offers the most developed exposition of this idea; for a

discussion of Bratman's view consonant with the one developed here see Butterfill 2015).

Though it has not received much attention in the literature on shared intentionality, an expressivist planning semantics has been developed for the deontic modalities. Gibbard (2003, chapters 3 and 4) and Stovall (Forthcoming B) offer the most precise formulations of this world-mind semantics, but it has its roots in a body of research Wilfrid Sellars began developing in the 1950s (e.g. 1951, 1963, 1966a, 1966b, 1967, chapter 7 of 1968, and 1976). The assertion of sentences falling under these modalities are interpreted as giving expression to plans of action concerning what one would do were one any member of a relevant group. In this regard deontic modality marks a capacity to put ourselves into different points of view within some community: the claim that cooks ought to clean their knives when dirty expresses commitment to the plan to clean one's knife in any situation in which one is a cook whose knife is dirty. Just so, features of the world that constrain or are otherwise conditions for the satisfaction of these plans can be accounted for in terms of possible worlds. The claim that my knife is dirty and I ought to wash it at once represents the world as being a particular way and expresses commitment to a plan to do something.

In order to account for the difference between the strong deontic modalities (ought and forbidden) and the weak (may and may not), we must distinguish two ways of choosing: *single-mindedly*, whereby one rejects all choices incompatible with the choice one makes, and *indifferently*, whereby the agent is capable of choosing single-mindedly and there is at least one other choice incompatible with the choice made which could have been made without changing any of the agent's single-minded choices (the capability to choose single-mindedly is a condition ensuring that non-human animals do not trivially count as choosing indifferently in the relevant sense). In choosing indifferently one has rejected rejecting the choices one could have taken without changing one's single-minded choices. A version of this distinction appears to have been independently developed at least three times over the last

half century, and it has not received the attention it deserves. Sellars marks it in at least two places (1966b: 113-5 and 1976: 57), though he does not discuss it in detail. Gibbard (2003 chapter 3) then proposes it, but without reference to Sellars. Finally Dreier (2006, 2009) advances a similar idea, also without referencing either Sellars or Gibbard. Koons (2019: 135) briefly discusses Sellars' account of indifference and Silk (2015: 53 fn. 8) notes that many of Gibbard's critics, by ignoring his discussion of this distinction, have not done justice to his view. Each of Sellars, Gibbard, and Dreier speaks of preference rather than single-mindedness as the contrast with indifference, but as any choice might be thought to exhibit one's preferences—and as the relevant notion is the doubled-stance one adopts both toward what one chooses and toward choices incompatible with what one chooses—I prefer to speak in terms of single-mindedness (furthermore, the definition of indifference in terms of single-mindedness is my own).

The distinction between these two ways of choosing, and its role in a planning semantics for deontic judgments, can be appreciated through an example. Suppose I am a cook who thinks I ought to wash my knife, that I may either wash it by hand or with the dishwasher, and that I ought not wash it with turpentine. I am thereby in a state of mind of planning on single-mindedly choosing to wash my knife (rejecting not washing it), single-mindedly choosing not to wash it with turpentine (rejecting washing it with turpentine), and indifferently choosing between washing it by hand or with the dishwasher (rejecting rejecting each of these choices). Even where I have a general policy of washing knives by hand, if I regard both that action and use of the dishwasher as permitted then I am planning on indifferently choosing to do the former and not the latter, in the sense introduced above: for I could just as well have chosen the latter without changing any of my single-minded choices. On the other hand, my choice to wash the knife by some means or other, as well as my choice not to use turpentine, are choices I undertake single-mindedly: for I could not make some

choice incompatible with either of these choices while remaining of the same mind. To choose indifferently among permitted actions each of which satisfies some single-minded choice is to discriminate *practical species* or determinations of a *practical genus* or determinable: e.g., washing a knife by hand and with the dishwasher are practical species of the practical genus of knife washing, and cognitive grasp of that relationship is reflected in the practical capacity to distinguish the corresponding single-minded and indifferent choices.

A semantics making use of these planning states offers a rigorous means for explicating the idea that descriptive sentences play a world-representing role while prescriptive sentences play an action-guiding or expressive role. Consider a language $\mathcal{L}$ composed of descriptive sentences, prescriptive sentences, and constructions from these and the Boolean operators (in the interest of readability I generally ignore use/mention distinctions below). Let a *possible world w* be defined as a maximally determinate state of affairs. That is, for every atomic descriptive sentence $d_{atom} \in \mathcal{L}$ and every w, either d_{atom} is true at w or $\sim d_{atom}$ is true at w (I will also speak of worlds as collections of facts or true propositions). 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* in *C* on h_D; or

2) α single-mindedly chooses not to *A* in *C* on h_D; or

3) α indifferently chooses to *A* in *C* on h_D; or

4) α indifferently chooses not to *A* in *C* on h_D.

Let a *deontic hyperstate* <w, h$_D$> be an ordered pair of a possible world and a deontic

hyperplan (I will occasionally speak simply in terms of worlds and plans).[5] With standard

definitions for the Boolean operators, any sentence of $\mathcal{L}$ can be uniformly interpreted in terms

of these hyperstates (see Forthcoming D for details).

This hyperstate semantics can be straightforwardly extended to account for individual

and shared intentions. We do so by supplementing deontic hyperplans with intentional

hyperplans (the need to keep two sets of books on planning will become clear in a moment).

An *intentional hyperplan* $h_I \in H_I$ 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 in C on h$_I$; or
2) α chooses not to A in C on h$_I$

Because the intentional modalities do not distinguish a strong and a weak modal force, the

distinction between single-mindedly and indifferently choosing to do something does not need

to be marked in the plans that model intentional states. In part V I will argue that this feature

of the semantics, together with the joint fact that 1) single-minded and indifferent choices

underlie the ability to mark off genus/species relations in practical reasoning, and 2) single-

minded and indifferent choice attitudes each require that an agent entertain attitudes toward

classes of action that are incompatible with what one chooses, suggests that the capacity for

shared intentionality is a more primitive form of the kind of practical rationality exhibited in

deontic cognition.

---

[5] I am treating the condition C that induces an action according to a plan as part of the planning content; it is also possible to model C with the world-content of hyperstates and treat intentions conditionally, e.g., "if in C, then we shall A".

A *deontic-intentional hyperstate* $\langle w, h_D, h_I \rangle \in S_{DI}$ consists of a set of ordered triples of a world, a deontic hyperplan, and an intentional hyperplan. Let s_{DI} range over elements of S_{DI}. We use deontic-intentional hyperstates to semantically evaluate (mental states associated with) the utterances of descriptive, prescriptive, and intentional sentences as follows. We denote the semantic interpretation of φ as $[[\varphi]]$. Descriptive sentences d are semantically evaluated as the worlds at which d is true:

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

Read this as 'the semantic value of the sentence d (under an interpretation) is the set of deontic-intentional hyperstates such that d is true at the world of those hyperstates'. The mental states associated with utterances of the three kinds of atomic prescriptive sentences are evaluated as follows:

"people are obliged to *A* in *C*" expresses universally rejecting not doing *A* in *C*

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

"people are forbidden to *A* in *C*" expresses universally rejecting doing *A* in *C*

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

"people are permitted to *A* in *C*" expresses universally rejecting rejecting doing *A* in *C*

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

Because negation is interpreted as a complement operator on sets of hyperstates, the usual equivalences among what is obliged, what is permitted, and what is forbidden hold.

To represent the contents of intentional sentences I use 'ι' as a metalinguistic variable taking either the first person singular or plural pronoun as its value. Formally, the semantic value of "ι shall A in C" is 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:

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

Here a concern might be raised that this does not deliver the right result as an analysis of shared intentionality. For there are cases where two people each want to do A in C, and want the other to do so, and yet do not share a plan. For instance, you and I might each have a plan to paint a house with the other, yet one of us plans to deceive the other in some of the details. Here we do not share the intention to paint the house together, or share it in the right way.

To avoid misclassifying such cases as instances where an intention is shared, in the literature on shared intentionality it is common to impose a requirement that the plans be executed in the right way. This may include Bratman's requirement that the plans 'interlock', Ludwig's requirement that each agent's plan include a specification of the plans of the other agents in the group, or Tuomela's requirement that we-mode shared intentions are acted on in part because of or for the reason that the intention is shared. In the interest of being as non-committal about these debates as I can, my official view is that, in cases of shared intentionality, the adverb 'together' is used in the specification of the action content of the plan (e.g., for a shared action A the action should be understood as *doing* A *together*), and I

assume that this stands as a placeholder for the conditions someone like Bratman, Ludwig, or Tuomela might impose. For present purposes I will also assume that whatever natural/scientific conditions satisfy *togetherness* in shared intentionality can serve here as well (cf. Butterfill 2018's use of entrainment, motor-representation mirroring, task co-representation, emergent coordination, collective goal-states and joint affordances as a basis for constructing a minimal theory of shared intentionality). So long as these conditions do not presuppose a capacity for planning single-mindedly, that assumption will not undercut the inferences about ontological priority made below.

We are now in a position to see how this plan-theoretic analysis of the deontic modalities supports an analogical understanding of shared intentionality. For this is a view of moral reasoning as a capacity to put ourselves into the perspectives of other people and to consider how we would behave were we in their position. That capacity may take effort to cultivate, as with the ability to enjoy the arts of a community we did not grow up in. But insofar as we possess this capacity, we represent our community as one whose members we identify with in the practical sense that we think about their actions as, after a fashion, actions that we ourselves take part in. The *moral* stance is the universalization of this transpersonal practical point of view, permuted through the category of single-mindedness. To put the point in Kantian terms, the universal kingdom of ends is one whose members settle on what morality requires of them independently of their person and place in space and time, and they pursue that kingdom single-mindedly. And on this semantics the transpersonal point of view is characteristic of shared intentionality as well.

IV A Commentary on the Model

As with any analogy, it is important to note where the similarity breaks down. In this section I provide some remarks on the single-minded/indifferent frame of mind and the merely

intentional, I discuss some of the ways they differ, and I show how these differences are represented in the semantics.

As I am using the notions, it is possible to choose without choosing either single-mindedly or indifferently. For the single-minded choice attitude is one that involves adopting an attitude of rejection toward all actions incompatible with what one chooses, and this requirement is not in force when simply choosing to do something (similar remarks hold for indifference). This is an idealization, of course, as we are at best approximate or partial hyperplanners. Much of what we do is habitual and unreflective, and when we do reflect on our actions we are often motivated by all sorts of things that we either are not aware of, or are actively trying to ignore. Still, our planning mental states can be modelled as approximations to what perhaps only the angels are capable of. Because single-mindedness and indifference are needed to model the strong and the weak deontic modalities, and because the intentional modalities lack a strong and a weak modal force, one can satisfy a merely intentional plan without choosing either single-mindedly or indifferently. At the same time, one way of making a choice is to choose single-mindedly (again, similarly for indifference). It follows that one way of realizing one's intentions is to do what one thinks one ought. That is, every single-minded choice to *A* is a choice to *A*, and so such a choice satisfies the intentional plan to *A*. But not every choice to *A* is single-minded, and so the mere fact that one chooses to *A* does not imply that one satisfies one's deontic plans—we may say that the agent did the right thing, but not for the right reason (or in the right frame of mind). This leaves open that the action may nevertheless be praiseworthy, as doing the right thing is often difficult enough on its own. But there is special praise that comes with doing the right thing under the right attitude (think of the difference between the executive forced by a court of law to donate to charity, and one who does so in the interest of the charity).

The fact that choosing simpliciter can be satisfied without an agent adopting any consideration toward incompatible choices, whereas the single-minded and indifferent choice attitudes cannot, is one reason for distinguishing intentional hyperplans from deontic hyperplans in the hyperstates that model the contents of theoretical and practical cognition, and so specifying where the analogy between these forms of thought breaks down. Another reason is provided by the following consideration. As exercises of practical rationality, in general (and where there are no defeaters) one who utters "I ought to A in C", just as one who utters "I shall A in C", will be disposed to A in C. But the two states of mind differ on the strength of the connection to the performance of the action in question, and we sublunary beings are characterized by the possibility of akrasia. For this reason, "I ought to A in C and I shall not A in C", though in some sense involving a kind of error in cognition, is not incoherent. This is a regrettable, but understandable, state of mind. The mental state associated with an assertion of this sentence is determined by the intersection of the two sets of deontic-intentional hyperstates determined by the two conjuncts, and owing to the fact that deontic and intentional sentences are evaluated relative to different sets of plans, there is a non-trivial semantic value for this mental state. If deontic and intentional cognition were alike evaluated according to the same sets of plans, however, this conjunction would be incoherent as there would be no set of hyperstates that model it. Nevertheless, there is a straightforward explanation for the practical irrationality of one who is in such a state: given the plans that constitute the contents of the separate conjuncts, there is no choice one can make that satisfies everything one has planned. Akrasia is therefore conceptually coherent, though practically irrational.

Note that this semantics makes no appeal to superminds in accounting for the contents of shared intentions. The only novel entities postulated are the plans that individual thinkers may have in mind. Concerning the subjects who have these plans in mind, we suppose that

they have the capacity, individually, to put themselves into the perspectives of other people and plan from these different points of view. That we are able to *collectively* have a plan in mind is a function of the fact that 1) plans are individuated relative to different agents and 2) we are able to consider how we would behave were we in various situations, including situations in which we are other people. Thus, numerically identical planning-contents can be shared across minds. But each point of evaluation at a plan is always the point of a *single individual*, so there is nothing in the semantics that commits us to anything but the mental states of individual people.

V A Hypothesis about Shared Intentionality and the Deontic Frame of Mind

At the beginning of the essay I noted that Hegel associated analogical inference with the same syllogistic figure that Peirce labelled 'hypothesis' and later assimilated to abduction. Both philosophers held that reasoning in this mode allows for changing the way we understand some target of inquiry, and they thought that a study of the operation of this reflective spontaneity (to think the idea in Kantian terms) was essential for future intellectual and social development. Just so, this analogical conception of the relationship between shared intentionality and deontic reasoning allows us to rethink our understanding of the relationship between individual and shared intentionality. In part II I suggested that some of the empirical work on shared intentionality and norm psychology supports my proposal; here I want to indicate how the empirical work might in turn be informed by this view. Despite not forcing our hand, the analysis does raise some intriguing hypotheses.

For instance, it falls out of my account that while moral reasoning and shared intentionality are of a common kind, the former involves a more fine-grained capacity for exercising practical rationality than is needed for the latter. To reason in a deontic frame of mind one must be able to choose single-mindedly, and the single-minded choice to do *A*

differs from merely *choosing* to do *A* insofar as it requires that one also bear a negative attitude toward actions that are incompatible with so choosing. Simply choosing to do *A*, however, amounts to stumping for *A*, and that choice can be made without regard to incompatible alternatives.[6] This conclusion is reinforced by a feature of practical rationality noted above. The ability to discriminate the indifferent choice from the single-minded choice is the ability to mark off the practical species or determinations that instantiate a practical genus or determinable. That is to say our *theoretically rational capacity* to make the deontic judgments we do is conditioned by our *practically rational capacity* to mark off genus/species or determinable/determination relations in agency. For an agent can only comprehend the requirement to do the right thing insofar as she is aware of what counts as 'the right thing', and in most of the contexts we face there are many ways of acting so as to satisfy this requirement. The world being the multifaceted thing it is, then, one can only plan from a deontic point of view insofar as one has the capacity to distinguish practical genera and their species, and this is conditioned by our ability to adopt single-minded and indifferent practical stances.

Finally, it is worth pointing out that a language of intentionality can be used to specify positive and negative reinforcement behaviors sufficient to habituate the members of a community into norm obedience of the sort expressed with the deontic modalities. And this can be done without presupposing that the individuals whose behaviors are shaped by this process have the capacity to discriminate single-mindedness from indifference—indeed, even the ones *instituting* this behavior might not be capable of explicit deontic reasoning. Where 'stickbeating' and 'flocktending' refers respectively to negative and positive behavioral

---

[6] For Hegel, this marks the difference between the merely desirous animal kingdom and the spiritual domain of human beings. When presented with food the mere animal will "fall to without ceremony and eat" (1977: 65—see the discussion in Brandom 2019: 240-3). We creatures of Spirit, on the other hand, are able to invest the world with value-laden significance, ceremonial or otherwise—food and drink are not merely ingested, but may become the body and blood of a divinity whose spirit suffuses the community collectively participating in this ceremony.

reinforcement, and understood on supposition that the hominids who were undergoing this conditioning shared a form of life (an affective/cognitive/practical stance toward a common environment), we have:[7]

- acting on the intention to stickbeat doing *A* in *C* and flocktend not doing *A* in *C* is one way of training people to conform to the norm that *A* is forbidden in *C*

- acting on the intention to stickbeat not doing *A* in *C* and flocktend doing *A* in *C* is one way of training people to conform to the norm that *A* is obliged in *C*

- acting on the intention to neither stickbeat doing *A* in *C* and nor flocktend not doing *A* in *C* is one way of training people to conform to the norm that *A* is permitted in *C*

It may be that the ability to *recognize* that something is obliged, permitted, or forbidden requires the ability to distinguish single-minded from indifferent choices. That is to say, self-consciousness may require the resources of deontic cognition. Nevertheless, the ability to *conform to the norms* governing what is obliged, permitted, and forbidden can be instituted within a community simply on the basis of behavior specifiable in terms of intentions unmarked by the distinction between single-mindedness and indifference. This again suggests that the merely intentional state of mind, whether individual or shared, is a more primitive form of cognition than that which attends deontic mental states. To return to Tomasello's philosophical anthropology, the hypothesis I am proposing—as a further specification of his account of the development of truth- and morality-tracking cognition from joint and then collective intentionality—is that a capacity for single-mindedness and

---

[7] The language of stickbeating is from Brandom (1994).

indifference had to emerge between the evolution of collective intentionality and the appearance of discursive theoretical and practical rationality. And I am suggesting that, nevertheless, shared intentional habit-shaping acculturating practices could have driven the development of a selection pressure in favor of that capacity among human beings, who could thereby more effectively control the development of the norms that bind them together.

VI Understanding Individual Intentionality by Analogy with Shared Intentionality

A second hypothesis raised by this framework concerns the status of shared intentions *vis à vis* individual intentions. It is generally held that individually intentional mental states are more basic than shared intentional mental states, with philosophers scaling up and supplementing the latter in the interest of explaining the former. By shifting to an analysis of the states of mind *expressed* with the language of the intentional modalities, and understanding these on the basis of a planning semantics for the deontic modalities, we arrive at a different view on the relationship between shared and individual intentions. For from the standpoint arrived at here it is more parsimonious to understand individual intentionality by analogy with shared intentionality rather than *vice versa*. Starting with the notion of a shared intentional state as a collection of hyperstates that are general with regard to the agents whose choices are planned, the mental state constituting an individual intention can be understood as the localization of that general set of plans to a single agent. Rather than seeing the shared intentional state of mind as an expansion of the individually intentional state of mind, we can see the individually intentional state of mind as a *contraction* of the shared intentional one. On this analysis, the universally practically-rational point of view is the conceptual foundation for the practically-rational mental states that express senses of self that are localized to particular communities (e.g. cooks), with the individual point of view as the limit of this circumscription of the primordial practical stance.

This hypothesis can be understood as one way of spelling out the idea that the shared intention is the locus of original intentionality within a community of rational beings. Regarding metaphysics or ontology, the only minds in play are those of the individual members of different communities. This is to say that the individual intention may be prior to the shared intention in the order of being (just as the shared intention, I have argued, is prior in being to deontic cognition). But in the order of understanding, as reflected in a hyperstate semantics for statements that give expression to intentions, it proves simpler to suppose that shared intentionality is basic and to define the mental state of an individual intention as a special case of the former (just as a planning semantics originally developed for the deontic modalities was positioned above as prior in the order of understanding to a planning semantics for intentionality). This line of thinking provides some support for the supposition that modern philosophy's interest in Spirit or the General Will was a conceptual advance over pre-modern notions of individual and shared agency. And, if the first hypothesis is correct, then this analogical investigation into the metaphysics of practical rationality helped discover something new about the nature of Spirit: *viz.*, that in the order of being the shared intention is prior to deontic cognition.

References

Bartha, Paul (2010) *By Parallel Reasoning: The Construction and Evaluation of Analogical Arguments*. New York: Oxford University Press.

Birch, Jonathan (2018). 'Joint Know-How.' *Philosophical Studies*, published online October 5, 2018, 1-24.

Brandom, R. B. (1994) *Making It Explicit: Reasoning, Representing, and Discursive Commitment*. Cambridge: Harvard University Press.

  --- (2019) *A Spirit of Trust: A Reading of Hegel's* Phenomenology. Cambridge: Harvard University Press.

Bratman, Michael E. (2014) *Shared Agency: A Planning Theory of Acting Together*. New York: Oxford University Press.

Burgh, Gilbert, Terri Field, and Mark Freakley (2006) *Ethics and the Community of Inquiry: Education for Deliberative Democracy*. South Melbourne: Thomson Social Science Press.

Butterfill, Stephen A. (2015) 'Planning for Collective Agency'. In Catrin Misselhorn (ed.), *Collective Agency and Cooperation in Natural and Artificial Systems: Philosophical Studies Series, Vol. 122* (Cham: Springer International), pp.149-168.

--- (2018) 'Coordinating Joint Action'. In Jankovic and Ludwig (2018), 68-82.

Chant, Sara Rachel (2018) 'Collective Action and Agency'. In Jankovic and Ludwig (2018), 13-24.

Chant, Sara Rachel, Frank Hindriks, and Gerhard Preyer (eds.) (2014) *From Individual to Collective Agency: New Essays*. New York: Oxford University Press.

Chudek, Maciej and Joseph Henrich (2011) 'Culture-Gene Coevolution, Norm-Psychology and the Emergence of Human Prosociality'. *Trends in Cognitive Sciences*, 15 (5), 218-226.

Dreier, James (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* (New York: Oxford University Press), pp. 217-233.

--- (2009) 'Relativism (and Expressivism) and the Problem of Disagreement,' *Philosophical Perspectives*, 23, 79-110.

Fletcher, Grace E., Felix Warneken, and Michael Tomasello (2012) 'Differences in Cognitive Processes Underlying the Collaborative Activities of Children and Chimpanzees'. *Cognitive Development*, 27, 136-153.

Gentner, Dedre (1983) 'Structure-Mapping: A Theoretical Framework for Analogy.' *Cognitive Science*, 7, 155-170.

Gentner, Dedre, Keith J. Holyoak and Boicho N. Kokinov (eds) (2001) *The Analogical Mind: Perspectives from Cognitive Science*. Cambridge: The MIT Press.

Gibbard, Allan (2003) *Thinking How to Live*. Cambridge: Harvard University Press.

Gilbert, Margaret (1989) *On Social Facts*. Princeton: Princeton University Press.

--- (2000) *Sociality and Responsibility: New Essays in Plural Subject Theory*. New York: Rowman and Littlefield.

--- (2014) *Joint Commitment: Howe We make the Social World*. New York: Oxford University Press.

Göckeritz, Susanne, Marco F.H. Schmidt, and Michael Tomasello (2014) 'Young Children's Creation and Transmission of Social Norms'. *Cognitive Development*, 30, 81-95.

Hegel, G. W. F. (1977) *Phenomenology of Spirit*. Translated by A.V. Miller, with analysis of the text and foreward by J. N. Findlay. Oxford: Oxford University Press.

Hesse, Mary (1952) 'Operational Definition and Analogy in Physical Theories.' *The British Journal for the Philosophy of Science*, 2 (8), 281-294.

--- (1966) *Models and Analogies in Science*. Notre Dame: University of Notre Dame Press.

Jankovic, Marija and Kirk Ludwig (eds.) (2018) *The Routledge Handbook of Collective Intentionality*. New York: Routledge.

Kant, Immanuel (1996). *Cambridge Edition of the Works of Immanuel Kant: Practical Philosophy*. Translated and edited by Mary J. Gregor. New York: Cambridge University Press.

Koons, Jeremy (2019) *The Ethics of Wilfrid Sellars*. New York: Routledge.

Koreň, Ladislav (2016) 'Joint Intentionality: From Thin to Thick'. *Journal of Social Ontology*, 2 (1), 75-85.

Ludwig, Kirk (2007) 'Collective Intentional Behavior from the Standpoint of Semantics'. *Noûs*, 41 (3), 355-393.

--- (2014) 'The Ontology of Collective Action'. In Chant, et al. (2014), 112-133.

--- (2016) *From Individual to Plural Agency: Collective Action: Volume 1*. New York: Oxford University Press.

--- (2018) 'Actions and Events in Plural Discourse'. In Jankovic and Ludwig (2018), 476-488.

Miller, Seumas (2010) *The Moral Foundations of Social Institutions: A Philosophical Study*. New York: Cambridge University Press.

--- (2014) 'Review of *Social Ontology: Collective Intentionality and Group Agents* by Raimo Tuomela'. *Notre Dame Philosophical Reviews*, April 20, 2014. https://ndpr.nd.edu/news/ocial-ontology-and-collective-intentionality-critical-essays-on-the-philosophy-of-raimo-tuomela-with-his-responses/

Miller, Kaarlo and Raimo Tuomela (2014) 'Collective Goals Analyzed'. In Chant et al. (2014), 34-60.

Pacherie, Elizabeth (2018) 'Collective Phenomenology'. In Jankovic and Ludwig (2018), 162-73.

Pettit, Philip (2003) 'Groups with Minds of their Own'. In Schmitt (2003a), 167-93.

Pettit, Philip and Christian List (2011) *Group Agency: The Possibility, Design and Status of Corporate Agents*. New York: Oxford University Press.

Price, Huw (2011) *Naturalism Without Mirrors*. New York: Oxford University Press.

--- (2013) *Expressivism, Pragmatism, and Representationalism*. New York: Cambridge University Press.

Rakoczy, Hannes, Felix Warneken, and Michael Tomasello (2008) 'The Sources of Normativity: Young Children's Awareness of the Normative Structure of Games'. *Developmental Psychology*, 44 (3), 875-881.

Redding, Paul (2003) 'Hegel and Peircean Abduction'. *European Journal of Philosophy*, 11 (3), 295-313.

Rizzolatti, Giacomo and Corrado Sinigaglia (2010) 'The Functional Role of the Parieto-Frontal Mirror Circuit: Interpretations and Misinterpretations'. *Nature Reviews Neuroscience*, 11 (4), 264-274.

Schmid, Hans Bernhard (2013) 'Shared Intentionality and the Origins of Human Communication'. In Alessandro Salice (ed.), *Intentionality* (München: Philosophia-Verlag), pp.349-368.

--- (2018) 'The Subject of 'We Intend''. *Phenomenology and the Cognitive Sciences*, 17 (2), 231-43.

Schmidt, Marco F.H. and Hannes Rakoczy (2018) 'Developing an Understanding of Normativity'. In Albert Newen, Leon De Bruin, and Shaun Gallagher (eds.), *The Oxford Handbook of 4E Cognition* (New York: Oxford University Press), pp.685-706.

Schmidt, Marco F.H., Hannes Rakoczy, and Michael Tomasello (2013) 'Young Children Understand and Defend the Entitlements of Others'. *Journal of Experimental Child Psychology*, 116 (4), 930-44.

Schmitt, Frederick F. (ed) (2003a) *Socializing Metaphysics: The Nature of Social Reality*. New York: Rowman and Littlefield.

--- (2003b) 'Joint Action: From Individualism to Supraindividualism'. In Schmitt (2003a), 129-65.

Schmitz, Michael, Beatrice Kobow, and Hans Bernhard Schmid (eds.) (2013) *The Background of Social Reality*. Dordrecht: Springer.

Schweikard, David P, and Hans Bernhard Schmid (2013) 'Collective Intentionality'. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy*. https://plato.stanford.edu/entries/collective-intentionality/.

Searle, John (1995) *The Construction of Social Reality*. New York: Free Press.

--- (2010) *Making the Social World: The Structure of Human Civilization*. New York: Oxford University Press.

Sellars, W. (1951) 'Obligation and Motivation'. *Philosophical Studies*, 2(2), 21-25.

--- (1963) 'Imperatives, Intentions, and the Logic of 'Ought''. In H-N. Casteñeda and G. Nakhnikian (eds.), *Morality and the Language of Conduct* (Detroit: Wayne State University Press), pp. 159-214.

--- (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 Keith Lehrer (ed), *Freedom and Determinism* (New York: Random House), pp.105-139.

--- (1967) '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.

--- (1968) *Science and Metaphysics: Variations on Kantian Themes*. Atascadero: Ridgeview Publishing Company (1992).

--- (1976) 'Volitions Re-Affirmed'. In M. Brand and D. Walton (eds.), *Action Theory* (Dordrecht: D. Reidel Publishing Company), pp.47-66.

Silk, Alex (2015) 'How to Be an Ethical Expressivist'. *Philosophy and Phenomenological Research*, XCI (1), 47-81.

Stovall, Preston (2015) 'Inference by Analogy and the Progress of Knowledge: From Reflection to Determination in Judgements of Natural Purpose'. *British Journal of the History of Philosophy*, 23 (4), 681-709.

--- (Forthcoming A) 'Syllogistic Reasoning as a Ground for the Content of Judgement: A Line of Thought from Kant through Hegel to Peirce'. *European Journal of Philosophy*.

--- (Forthcoming B) 'Understanding What We Ought and Shall Do: A Hyperstate Semantics for Descriptive, Prescriptive, and Intentional Sentences'. In Ladislav Koreň, Hans Bernhard Schmid, Preston Stovall, and Leo Townsend (eds.), *Groups, Norms, and Practices* (Dordrecht: Springer).

--- (Forthcoming C) 'Normative Attitudes, Collective Intentionality, and Discursive Cognition'. In Leo Townsend, Hans Bernhard Schmid, and Preston Stovall (eds.), *The Social Institution of Discursive Norms* (New York: Routledge).

--- (Forthcoming D) *The Single-Minded Animal: Shared Intentionality, Normativity, and the Foundations of Discursive Cognition*. New York: Routledge.

Tomasello, Michael, Malinda Carpenter, Josep Call, Tanya Behne, and Henrike Moll (2005) 'Understanding and Sharing Intentions: The Origins of Cultural Cognition'. *Behavioral and Brain Sciences*, 28 (5), 675-91.

Tomasello, Michael (2014) *A Natural History of Human Thinking*. Cambridge: Harvard University Press.

Tuomela, Raimo (2007) *The Philosophy of Sociality: The Shared Point of View*, New York: Oxford University Press.

--- (2013) *Social Ontology: Collective Intentionality and Group Agents*. New York: Oxford University Press.

Tuomela, Raimo and Kaarlo Miller (1988) 'We-Intentions'. *Philosophical Studies*, 53, 367-389.