

Reasoning with Moral Conflicts

JOHN F. HORTY

University of Maryland, College Park

1 Introduction

Let us say that a *normative conflict* is a situation in which an agent ought to perform an action *A*, and also ought to perform an action *B*, but in which it is impossible for the agent to perform both *A* and *B*. Not all normative conflicts are moral conflicts, of course. It may be that the agent ought to perform the action *A* for reasons of personal generosity, but ought to perform the action *B* for reasons of prudence: perhaps *A* involves buying a lavish gift for a friend, while *B* involves depositing a certain amount of money in the bank. In general, our practical deliberation is shaped by a concern with a variety of morally neutral goods—not just generosity and prudence, but any number of others, such as etiquette, aesthetics, fun—many of which are capable of providing conflicting reasons for action. I mention these ancillary values in the present setting, however, only to put them aside. We will be concerned here, not with normative conflicts more generally, but precisely with *moral conflicts*—situations in which, even when our attention is restricted entirely to moral reasons for action, it is nevertheless true that an agent ought to do *A* and ought to do *B*, where it is impossible to do both.

It is often argued that moral conflicts, defined in this way, simply cannot occur, that they are impossible. The justifications offered for this conclusion fall into two broad categories. Some writers contend that, although there might be normative conflicts more generally, the possibility of specifically moral conflicts is ruled out by the special nature of moral reasons. Arguments along these lines generally proceed by identifying as genuinely moral reasons for action only those supported by some particular moral theory—usually a Kantian or utilitarian theory—that itself rules out the possibility of conflicts. Alan Donagan, for example, argues against moral conflicts in his [1984] and [1993] by advancing a kind of rationalist theory, developed

through a process of dialectical reasoning, according to which it is very nearly analytic that such conflicts cannot arise: whenever an apparent conflict is discovered, this is supposed to show only that the theory as developed thus far is defective, requiring further revision until the conflict is avoided. And of course, it is most natural also for an advocate of the utilitarian approach to be drawn toward Mill's own conclusion that the principle of utility, the ultimate moral reason, provides a common standard through which any apparent moral conflicts can be resolved.¹

I will have very little to say about this first style of argument, which denies the possibility of moral conflicts by appealing to considerations concerning the kinds of reasons for action that might be supplied by the correct moral theory; the general line of reasoning is sensible, of course, but the project of developing any such argument in detail would be a substantial task, since it requires the defense of some particular moral theory as correct. My attention here will be concentrated instead on a different style of argument, which denies the possibility of moral conflict, not so much by appealing to a particular moral theory as correct, but rather on the basis of broader conceptual considerations, sometimes, but not always, involving issues in deontic logic.

Generally, those who argue in this way—including Philippa Foot [1983], John Searle [1980], Judith Jarvis Thomson [1990], and more recently, David Brink [1994] and Paul Pietroski [1993]—are careful to distinguish between two different kinds of ought, or obligation, statements. Although the exact character of this distinction varies from one writer to the next, the basic idea is plain. There are supposed to be, first of all, statements describing broad moral reasons for action, which it is useful to think of as imperatives issued by some source of moral authority, or value. Since an agent might recognize different sources of value, and since even the same source of value can at times issue inconsistent imperatives, it is generally acknowledged that ought statements of this first kind might conflict. Adapting the well-known example from Sartre's [1946], we can imagine an agent whose conception of patriotism leads him to accept the imperative "Given the need, you ought to fight for your country," but whose conception of personal devotion leads him to accept the imperative "Given the need, you ought to care for an aging relative." The agent would then confront conflicting moral reasons for action in any situation in which he is needed both to fight for his country and to care for an aging relative, but can do only one or the other.

In addition to this first kind of ought statement—representing only moral reasons for action, and allowing for conflicts among these reasons—there is also supposed to be a second, distinct kind of ought statement, describing what one ought to do in a particular situation once all the relevant moral reasons are taken into consideration and weighed against each other. Ought statements of this second kind are, in a sense, derived from the first, since they are based on the variety of moral reasons bearing on a given situation.

But since they reflect the result of integrating and balancing these various reasons, it is thought that there can be no conflicts among ought statements of this second kind—that we cannot accept both the statement “Under the circumstances, you ought all things considered to defend your country,” as well as the statement the statement “Under the circumstances, you ought all things considered to care for your relative.”

We can mark the difference between these two kinds of ought statements—or more simply, *oughts*—by referring to the first as *prima facie* and to the second as *all things considered* oughts, although again, different writers employ different terminology to characterize the precise distinctions they have in mind. Using this language, my concern in the present paper is with the claim that, although there may be conflicts among *prima facie* oughts, there can be no moral conflicts involving all things considered oughts; and I focus special attention on a recent proposal—hinted at by Donagan and Foot, explicitly defended by Brink—known as the “disjunctive account.” The strongest case for moral conflicts seems to arise in situations in which the *prima facie* reasons for performing each of two incompatible actions, *A* and *B*, are either evenly balanced or else incommensurable. According to the disjunctive account, the correct all things considered conclusion to draw in these situations is, not that the agent ought to perform the action *A* and ought also to perform the action *B*, but simply that the agent ought to perform either *A* or *B*. In Sartre’s case, for example, the disjunctive account would lead us to conclude, not that the agent ought to defend his country and also that he ought to care for his relative, but simply that the agent ought either to defend his country or care for his relative, that he cannot neglect both duties.

The remainder of the paper is organized as follows. In the next section, I set out two very simple and closely related deontic logics, both designed for deriving all things considered oughts as conclusions from a set of *prima facie* oughts, possibly conflicting, taken as premises. One of these logics, although itself consistent, allows moral conflicts among the derived all things considered oughts. The other avoids moral conflicts by adopting the disjunctive account—providing, as far as I know, the first accurate formulation of this view. These two logics are not presented here for the sake of any particular technical interest; indeed, the presentation is rudimentary, and any unnecessary technical development is avoided. Instead, the point of the two logics is simply to provide a concrete illustration of two different strategies for reasoning in the face of conflicting *prima facie* oughts, as well as a clear conceptual framework within which issues involving the acceptability of all things considered conflicts can be addressed with some degree of precision. The following section is then devoted to an examination, within this framework, of some of the most important recent arguments on the topic. My conclusion is that, given the terms of the current discussion—that is, without appealing to any constraint on the

structure of moral reasons that might be provided by some particular moral theory—there is no logical or conceptual reason to reject the possibility of moral conflict.

2 The two logics

We will assume as background an ordinary propositional logic, containing the usual connectives, with the turnstile \vdash representing ordinary logical consequence. Purely for the sake of convenience, in order to avoid too much awkwardness in our formalization of particular examples, we will suppose that the background language allows for materially inconsistent atomic formulas—representing statements, like “It’s summer” and “It’s winter,” that cannot both be true at once even though neither is explicitly represented as a negation of the other.

This background logic is then supplemented with two different deontic operators corresponding to the two different kinds of oughts under consideration. Where A and B are statements from the background language, we let the formula $!(B/A)$ represent the idea that, under the circumstances A , it ought prima facie to be the case that B . The more conventional deontic formula $\bigcirc(B/A)$ will be reserved to express the idea that, under the circumstances A , it ought all things considered to be the case that B .

The representation of prima facie oughts is meant to suggest a picture of these statements as conditional imperatives, and it will be convenient to introduce two functions—*Antecedent* and *Consequent*—allowing us to pick out their antecedent and consequent parts: if i represents the prima facie ought $!(B/A)$, for example, then *Antecedent* $[i]$ is the statement A and *Consequent* $[i]$ is the statement B . The antecedent of a prima facie ought is a sort of triggering condition, specifying the circumstances under which that ought is relevant. The consequent of the ought determines its satisfaction conditions; we will say that a particular prima facie ought is *satisfied* whenever its consequent is true.

Our notation calls for two immediate comments. First, although the two kinds of oughts introduced here are both conditional, it is easy enough, as well as standard practice, to define their unconditional analogues as conditional oughts that happen to be conditioned only on the special proposition \top , representing a tautology, and so true in any situation whatsoever. The statement $!(B)$, meaning simply that it ought prima facie to be the case that B , can thus be taken as an abbreviation of the formula $!(B/\top)$, and the statement $\bigcirc(B)$, meaning that it ought all things considered to be the case that B , can likewise be taken to abbreviate the formula $\bigcirc(B/\top)$.

Second, the reader may have noticed that we have shifted from an informal discussion largely focused on questions concerning what various agents ought to do to a formal notation containing statements only about what ought to be the case. Although it is often important to distinguish personal

from impersonal oughts—statements about what agents ought to do, from statements about what ought to be the case—I believe that the issues raised by that distinction are orthogonal to the problems considered here: prima facie conflicts can arise concerning either personal or impersonal oughts, and a strategy for handling conflicts of either kind should be applicable also to the other. The present paper therefore follows a policy of intentional but, I hope, harmless equivocation. We will generally rely on personal oughts in our informal discussion, for the simple reason that they allow for the formulation of somewhat more natural examples; but in order to avoid extraneous complications involving the proper treatment of personal agency, which would be necessary for a full logical representation of these examples, the formal development itself will be restricted to the simpler case of impersonal oughts.²

With these preliminaries behind us, we can now turn to the general problem at issue in this paper: given a background context including an arbitrary set of prima facie oughts, how do we determine whether a particular all things considered ought holds under some specified set of circumstances? This general problem can be cast as a logical question: where the circumstances under consideration are specified by the formula A , how do we define a consequence relation determining whether a particular all things considered ought of the form $\bigcirc(B/A)$ —telling us that B ought to be the case under the circumstances A —follows from a context of prima facie oughts?

In the present paper, this question is answered in two steps, which we take up in turn. Not every prima facie ought needs to be satisfied in every situation, of course: a conditional imperative telling me how to act once I have made a promise may have no bearing whatsoever on a situation in which I have received an unsolicited request for a charitable donation, for example. The first step in determining whether an all things considered ought of the form $\bigcirc(B/A)$ follows as a consequence from a background context of prima facie oughts, then, is to identify the particular prima facie oughts from the context that are relevant under the circumstances specified by A , those that must be satisfied; these prima facie oughts can be described as *binding*. Once we have identified the set of prima facie oughts that are to be classified as binding under the circumstances specified by A , the second step is to describe the way in which the enjoined formula B is to be calculated from these binding prima facie oughts.

2.1 *Binding oughts*

Let \mathcal{I} represent the entire set of prima facie oughts from the background context. As an initial suggestion, it might seem reasonable to identify the oughts from \mathcal{I} that are to be classified as binding in particular circumstances with those whose antecedents are guaranteed to hold under those circumstances. Focusing on the role of antecedents as triggering conditions, we can

describe these prima facie as those that are *triggered* under the circumstances; and we can refer to the entire set of prima facie oughts that are triggered under the circumstances specified by A as $Triggered_{\mathcal{I}}(A)$, where this notion is defined as follows.

- Let \mathcal{I} be a set of prima facie oughts. Then the set of oughts from \mathcal{I} that are triggered under the circumstances A is

$$Triggered_{\mathcal{I}}(A) = \{i \in \mathcal{I} : A \vdash \textit{Antecedent}[i]\}.$$

It turns out, however, that this initial suggestion—that the binding oughts can simply be identified with the triggered oughts—is too liberal, forcing us at times to classify too many prima facie oughts as binding. The point can be illustrated through a standard example.

Imagine that my background set \mathcal{I} of prima facie oughts contains exactly two imperatives: “I ought to meet a friend for lunch, given that I have promised to do so” and “I ought to rescue a drowning child, given the need.” If we take the statement letters P , M , N , and R to stand for the respective propositions that I promise to meet a friend for lunch, that I actually meet my friend, that I am needed to rescue a drowning child, and that I actually rescue the child, these two prima facie oughts can then be represented as

$$\begin{aligned} i_1 &= !(M/P), \\ i_2 &= !(R/N). \end{aligned}$$

Now suppose, as the example goes, that I find myself in circumstances satisfying the condition $P \wedge N$, in which I have promised to meet a friend for lunch and am also needed to rescue a drowning child, but in which it is impossible for me both to meet my friend and also to carry out the rescue: the statements M and R are materially inconsistent. Under these circumstances, both of the two prima facie oughts i_1 and i_2 would be triggered, since the antecedent of each is entailed by the description of the situation: we would have $Triggered_{\mathcal{I}}(P \wedge N) = \{i_1, i_2\}$. According to our initial suggestion, then, which simply identifies binding oughts with triggered oughts, both of these two prima facie oughts would also have to be classified as binding. But if binding oughts are those that must be satisfied, this result seems to be wrong, from an intuitive point of view. It seems to be much more natural to classify only the second of these two prima facie oughts—rescuing the child—as binding, since they cannot both be satisfied, and since the second is so much more important.

In order to provide a precise characterization of judgments of importance like this, we therefore supplement our general framework with a preference ordering \leq , representing the relative priority among various prima facie

oughts. As the notation suggests, the basic ordering \leq encodes only a weak preference comparison: where i and j are two *prima facie* oughts, the statement $i \leq j$ is taken to mean that j is at least as important as i . The weak preference ordering \leq can be used, however, to define a corresponding relation $<$ of strict preference, with the statement $i < j$ —defined as true whenever $i \leq j$ and it is not the case that $j \leq i$ —taken to mean that j is strictly more important than i . Returning to our example, the strict preference for rescuing a child over meeting a friend for lunch can then be captured by stipulating that $i_1 < i_2$ —that is, $i_1 \leq i_2$ and it is not the case that $i_2 \leq i_1$.

What properties should we expect to find in these preference relations among *prima facie* oughts? It is reasonable, first of all, to assume that the weak preference relation \leq should satisfy the *reflexivity* property

$$i \leq i,$$

according to which any *prima facie* ought i is at least as important as itself; and it seems equally reasonable to assume the *transitivity* property

$$i \leq j \text{ and } j \leq k \text{ imply } i \leq k,$$

according to which the *prima facie* ought k is at least as important as i whenever k is at least as important as the *prima facie* ought j , and j itself is at least as important as i . A weak ordering relation \leq satisfying both reflexivity and transitivity is known as a *quasi-ordering*. The corresponding strong ordering, with $<$ defined as above, also satisfies transitivity, but fails reflexivity.

Should we assume any other properties in the preference orderings? In particular, should we assume that this ordering satisfies the property of *strong connectivity*, according to which

$$i < j \text{ or } j < i$$

holds whenever i and j are distinct *prima facie* oughts? Here we reach an important branch point in our discussion. On one hand, this strong connectivity assumption would allow for a convenient resolution to any potential moral conflict. What the assumption tells us is that, of any two *prima facie* oughts, one is always strictly more important than the other; and so it would be natural, in case of a conflict, to settle the matter simply by favoring the more important of the two.

On the other hand, the strong connectivity assumption is not particularly plausible on the face of it, and for two reasons. First, some pairs of *prima facie* oughts might seem to be incommensurable in value, as illustrated by Sartre's example, mentioned earlier, of a conflict involving the imperatives to defend one's country and to care for an aging relative. In this case, it

could be argued that the two imperatives involved issue from entirely separate sources of value—duty to country, versus duty to family—and cannot meaningfully be compared in importance. Second, even if two *prima facie* oughts do happen to issue from the same source of value, and are therefore comparable in importance, they might nevertheless violate the strong connectivity assumption, according to which one or the other must always be strictly more important, simply by having equal importance. As an example, suppose I have simultaneously arranged to have a private dinner this evening with each of two identical and identically situated twins, both of whom would now be equally disappointed by my cancellation; the situation can be made arbitrarily symmetrical.³ The resulting *prima facie* oughts—to have dinner with one twin, and to have dinner with the other—issue from the same source of value, and can meaningfully be compared in importance. But in light of the symmetry, what reason could there be for preferring one over the other?

Although they did not use the technical language of ordering or connectivity, some historical figures—Bradley, several of the British intuitionists, such as Ross—did seem to feel that *prima facie* oughts, or moral imperatives, could always be compared in importance and ranked in such a way that any potential moral conflicts would be resolved, if not abstractly, then at least in their application to a particular situation. However, both the process through which such a ranking could be arrived at and the grounds on which it might be defended have remained somewhat mysterious. Notoriously, Bradley and Ross, both influenced by Aristotle, imagined that the relative importance of the various *prima facie* oughts bearing on a particular situation could be discovered, and perhaps justified, simply through an intuitive appraisal—a kind of perceptual judgment made by the practically wise person, or in the case of Bradley, by the person who has properly identified his will with the moral spirit of the community. Of course, other writers in the pluralist tradition, also working with *prima facie* oughts deriving from separate sources of value, have attempted to describe a more theoretically transparent, and rationally defensible, procedure through which conflicts among these *prima facie* oughts might be adjudicated. However, although I do not know of any general argument against this possibility, I do think it is fair to say that all of the various procedures that have been elaborated to date either fail to guarantee that the conflicts will actually be resolved, or else rely, at some point, on a kind of moral insight no less obscure than that suggested by Bradley and Ross.⁴ In light of the apparent counterexamples to the strong connectivity assumption, then, and lacking any real justification for the idea, we will therefore suppose throughout the remainder of this paper that the ordering relation on *prima facie* oughts satisfies only the two quasi-ordering constraints, reflexivity and transitivity, allowing for the possibility of conflicting oughts that are either incomparable or identical in importance.

We can now return to the task of identifying the prima facie oughts that should be classified as binding in a particular set of circumstances—this time, however, explicitly taking as background, not simply a set \mathcal{I} of prima facie oughts, but a slightly more complicated structure of the form $\langle \mathcal{I}, \leq \rangle$, in which \leq is a quasi-ordering reflecting the relative importance of the various prima facie oughts belonging to \mathcal{I} . Such a structure will be referred to as a *background context* of prima facie oughts, or simply as a *context*; the idea is that it is a structure of this form, a quasi-ordered set of prima facie oughts, from which an agent's all things considered moral judgments are to be derived.

Our initial suggestion—that the binding oughts should be identified with the triggered oughts—was seen to be problematic in cases in which a triggered ought happens to conflict with a more important triggered ought. The example illustrating this problem, involving a conflict between meeting a friend and rescuing a child, can now be represented through the context $\langle \mathcal{I}, \leq \rangle$, where the set \mathcal{I} contains the two prima facie oughts i_1 and i_2 , telling me respectively to meet a friend if promised and to rescue a child if needed, and where the ordering relation \leq is defined so that $i_1 < i_2$, reflecting our preference for rescuing a child over meeting a friend. In this case, we saw that both of the two prima facie oughts i_1 and i_2 are triggered under the circumstances $P \wedge N$, in which I have promised to meet a friend but am also needed to rescue a child, but that, of these two triggered oughts, i_1 is not to be classified as binding since it conflicts with the more important triggered ought i_2 .

Generalizing from this example, let us say that a prima facie ought, even if it is triggered in some situation, is *defeated* whenever it happens to conflict with a more important prima facie ought that is also triggered in that situation, in the sense that the consequents of the two oughts are inconsistent.⁵ We can then arrive at a final characterization of the binding oughts through a slight modification of our initial suggestion, defining the set of prima facie oughts from a background context $\langle \mathcal{I}, \leq \rangle$ that are to be classified as binding under the circumstances A —represented as $Binding_{\langle \mathcal{I}, \leq \rangle}(A)$ —are those that are triggered, but also not defeated. The idea can be captured formally as follows.

- Let $\langle \mathcal{I}, \leq \rangle$ be a background context of prima facie oughts. Then the set of oughts from \mathcal{I} that are to be classified as binding under the circumstances A is

$$Binding_{\langle \mathcal{I}, \leq \rangle}(A) = \{i \in \mathcal{I} : (1) i \in Triggered_{\mathcal{I}}(A), \\ (2) \text{there is no } j \in Triggered_{\mathcal{I}}(A) \text{ such that} \\ (a) i < j, \\ (b) Consequent[i] \text{ and } Consequent[j] \\ \text{are inconsistent}\}.$$

Evidently, part (1) of this definition tells us that a binding ought must be triggered, and part (2) that it cannot be defeated: there can be no other

triggered ought that is both (a) more important and (b) conflicting. Applied to our example, we can now see that $Binding_{\langle \mathcal{I}, \leq \rangle}(P \wedge N) = \{i_2\}$, as desired. The prima facie ought i_2 is correctly classified as binding under the condition $P \wedge N$, since it is triggered but not defeated; but the prima facie ought i_1 is defeated, since i_2 is triggered, the preference ordering tells us that $i_1 < i_2$, and $Consequent[i_1]$ and $Consequent[i_2]$ are inconsistent.

2.2 Defining the logics

Having identified the prima facie oughts that are to be classified as binding under a certain set of circumstances as those that are triggered but not defeated, we now turn to the second step in the process of defining our two logics: determining how the all things considered oughts are to be calculated from the binding prima facie oughts.

Since the binding prima facie oughts are those that must be satisfied, the obvious idea would be to define the all things considered oughts as those that result when all the binding oughts are in fact satisfied—endorsing a statement of the form $\bigcirc(B/A)$, that is, whenever B is a necessary condition for satisfying all the prima facie oughts belonging to $Binding_{\langle \mathcal{I}, \leq \rangle}(A)$, the set of oughts classified as binding under the circumstances specified by the statement A .

In order to express this idea precisely, we first generalize the function *Consequent* so as to apply, not only to individual prima facie oughts, but also to sets of these oughts: where \mathcal{F} is a set of prima facie oughts, we now take

$$Consequent[\mathcal{F}] = \{Consequent[i] : i \in \mathcal{F}\}$$

to be the set containing the consequents of the various oughts belonging to \mathcal{F} . Since $Binding_{\langle \mathcal{I}, \leq \rangle}(A)$ is the set of prima facie oughts that are classified as binding under the circumstances A , the set of consequents of these oughts is $Consequent[Binding_{\langle \mathcal{I}, \leq \rangle}(A)]$; and so the binding oughts are satisfied whenever all the various statements belonging to the set $Consequent[Binding_{\langle \mathcal{I}, \leq \rangle}(A)]$ are true. The obvious idea—that the all things considered oughts are those that result from satisfying all the binding prima facie oughts—can therefore be formulated as follows: given a background context $\langle \mathcal{I}, \leq \rangle$, we should accept an all things considered ought of the form $\bigcirc(B/A)$ as a consequence whenever the statement B follows as an ordinary logical consequence from the statement set $Consequent[Binding_{\langle \mathcal{I}, \leq \rangle}(A)]$.

To illustrate, in the case of our earlier context $\langle \mathcal{I}, \leq \rangle$, where $\mathcal{I} = \{i_1, i_2\}$ and $i_1 < i_2$, we saw that the set of prima facie oughts that are binding under the condition $P \wedge N$ is $Binding_{\langle \mathcal{I}, \leq \rangle}(P \wedge N) = \{i_2\}$. The set of consequents of these binding oughts is therefore $Consequent[Binding_{\langle \mathcal{I}, \leq \rangle}(P \wedge N)] = \{R\}$. And since the statement R is itself, of course, logically entailed by this set, the proposal would yield the all things considered ought $\bigcirc(R/P \wedge N)$ as a consequence of the background context—telling us, correctly, that what I ought to do, under the condition that I have promised to meet a friend

for lunch but am also needed to rescue a drowning child, is rescue the child.

In fact, this obvious proposal would be entirely adequate if the set of binding prima facie oughts were guaranteed to be conflict free—that is, if $\text{Consequent}[\text{Binding}_{\langle \mathcal{I}, \leq \rangle}(A)]$ were guaranteed to be consistent under any circumstances A . And it is easy to see that this set would have to be consistent if we had been able to adopt the assumption of a strongly connected preference ordering on prima facie oughts. But without this assumption—that is, given that we allow prima facie oughts to be either incommensurable or equal in importance—there is no guarantee that the set of binding prima facie oughts will be conflict free, and the obvious proposal leads to problems.

As an example, let us return to the situation in which I have inadvertently arranged to have a private dinner tonight with each of two twins, so that I am faced with conflicting but equally important prima facie oughts. Suppose that A_1 and A_2 stand for the respective statements that I have arranged to have dinner with twins 1 and 2, and that D_1 and D_2 stand for the statements that I will in fact have dinner with twins 1 and 2, where, since I cannot have a dinner with both, the statement set $\{D_1, D_2\}$ is inconsistent. In this case, the background context is $\langle \mathcal{I}, \leq \rangle$, where the two prima facie oughts belonging to \mathcal{I} are

$$\begin{aligned} i_3 &= !(D_1/A_1), \\ i_4 &= !(D_2/A_2), \end{aligned}$$

telling me that I ought to have dinner with each twin given that I have arranged to do so, and where the preference ordering holds that each of these oughts is at least as important as the other—that both $i_3 \leq i_4$ and $i_4 \leq i_3$ —so that neither is strictly more important. Under the condition $A_1 \wedge A_2$, then, where I have arranged to have dinner with both twins, both of these prima facie oughts are triggered and neither is defeated. We therefore have $\text{Binding}_{\langle \mathcal{I}, \leq \rangle}(A_1 \wedge A_2) = \{i_3, i_4\}$ as the set of binding prima facie oughts, and so $\text{Consequent}[\text{Binding}_{\langle \mathcal{I}, \leq \rangle}(A_1 \wedge A_2)] = \{D_1, D_2\}$ as the set of their consequents. Since this set is inconsistent, it entails any statement at all, of course, and so the obvious proposal, as sketched above, would force us to accept an all things considered ought of the form $\bigcirc(B/A_1 \wedge A_2)$ for any statement B whatsoever. But this is surely incorrect. Even if I have run into a sort of local difficulty by overbooking my evening, it would be odd to conclude from this that I ought to do absolutely everything.

Given that we cannot, therefore, adopt the obvious idea of defining the all things considered oughts as those that result from satisfying the entire set of binding prima facie oughts, what other options are available? Well, if we cannot satisfy all the binding prima facie oughts, it seems reasonable to satisfy as large a subset of them as we possibly can, at least—some subset of

the binding prima facie oughts that is itself satisfiable, but large enough that supplementing it with even one binding prima facie ought that it does not already contain would render it unsatisfiable. In order to develop a proposal along these lines, we first introduce the familiar notion of a maximal consistent subset of a set of formulas.

- Where \mathcal{H} is a set of formulas, \mathcal{F} is a *maximal consistent subset* of \mathcal{H} just in case (i) $\mathcal{F} \subseteq \mathcal{H}$, (ii) \mathcal{F} is consistent, and (iii) there is no consistent set \mathcal{G} such that $\mathcal{F} \subset \mathcal{G}$ and $\mathcal{G} \subseteq \mathcal{H}$.

In this definition, clauses (i) and (ii) tell us that \mathcal{F} is both a subset of \mathcal{H} and itself consistent; clause (iii) tells us that \mathcal{F} is as large a consistent subset of \mathcal{H} as possible, in the sense that supplementing it with even one additional element from \mathcal{H} would result in an inconsistent set \mathcal{G} . Using this notion of maximal consistency, then, the general strategy under consideration is to try to define the conditions under which we would wish to endorse an all things considered ought of the form $\bigcirc(B/A)$ by focusing, not on the entire set $\text{Consequent}[\text{Binding}_{\langle \mathcal{I}, \leq \rangle}(A)]$, containing the consequents of all the binding prima facie oughts, but only on the maximal consistent subsets of this set.

Unfortunately, the decision to follow this general strategy does not yet determine a unique approach, since it is possible for an inconsistent set of statements to contain more than one maximal consistent subset. As a result, there are two natural ways in which the general strategy of focusing on maximal consistent subsets could be developed. We might decide, as a first alternative, to endorse those all things considered oughts that result from satisfying any one of the various maximal consistent subset of the binding prima facie oughts—any of the various subsets of the binding prima facie oughts, that is, whose consequents form a maximal consistent subset of the entire set of consequents. This alternative will be described here as the *conflict account*, since it allows for conflicts among all things considered oughts.

- Let $\langle \mathcal{I}, \leq \rangle$ be a background context of prima facie oughts. Then the all things considered ought $\bigcirc(B/A)$ follows as a consequence of the context $\langle \mathcal{I}, \leq \rangle$ according to the *conflict account* if and only if $\mathcal{F} \vdash B$ for some maximal consistent subset \mathcal{F} of the set $\text{Consequent}[\text{Binding}_{\langle \mathcal{I}, \leq \rangle}(A)]$.

Broadly speaking, the idea underlying the conflict account is that a conclusion can be drawn from a body of information that may be inconsistent whenever one coherent way of looking at things—in this case, one maximal consistent subset of the binding prima facie oughts—supports that conclusion. An alternative idea is that a conclusion can be drawn only when it is supported by each way of looking at things, each maximal consistent subset of the binding oughts. This alternative yields the *disjunctive account*;

although it differs from the conflict account only in a single word—“some” is replaced by “each”—we set it out here in full for the sake of completeness.

- Let $\langle \mathcal{I}, \leq \rangle$ be a background context of prima facie oughts. Then the all things considered ought $\bigcirc(B/A)$ follows as a consequence of the context $\langle \mathcal{I}, \leq \rangle$ according to the *disjunctive account* if and only if $\mathcal{F} \vdash B$ for each maximal consistent subset \mathcal{F} of the set $\text{Consequent}[\text{Binding}_{\langle \mathcal{I}, \leq \rangle}(A)]$.

The differences between these two accounts can be illustrated by returning to the twin example, where, as we have seen, $\text{Binding}_{\langle \mathcal{I}, \leq \rangle}(A_1 \wedge A_2) = \{i_3, i_4\}$ is the set of binding oughts in the situation in which I have arranged to have dinner with both twins, so that the set of consequents of these oughts is the inconsistent set $\text{Consequent}[\text{Binding}_{\langle \mathcal{I}, \leq \rangle}(A_1 \wedge A_2)] = \{D_1, D_2\}$. Evidently, this inconsistent set of consequents has two maximal consistent subsets, $\mathcal{F}_1 = \{D_1\}$ and $\mathcal{F}_2 = \{D_2\}$. The conflict account thus supports both $\bigcirc(D_1/A_1 \wedge A_2)$ and $\bigcirc(D_2/A_1 \wedge A_2)$ as consequences, since \mathcal{F}_1 logically entails D_1 and \mathcal{F}_2 logically entails D_2 . The result is a conflict among all things considered oughts, telling me that I ought to have dinner with twin 1, and also that I ought to have dinner with twin 2, though I cannot do both. According to the disjunctive account, on the other hand, neither $\bigcirc(D_1/A_1 \wedge A_2)$ nor $\bigcirc(D_2/A_1 \wedge A_2)$ is supported, since neither D_1 nor D_2 is entailed by both of the two maximal consistent subsets \mathcal{F}_1 and \mathcal{F}_2 ; but of course, both \mathcal{F}_1 and \mathcal{F}_2 do entail the statement $D_1 \vee D_2$, and so the disjunctive all things considered ought $\bigcirc(D_1 \vee D_2/A_1 \wedge A_2)$ is supported. In the case of this example, then, rather than telling me, if I have arranged to dine with each twin but cannot in fact dine with both, that I nevertheless ought to dine with both and so face a moral conflict, the disjunctive account tells me only that what I ought to do, all things considered, is dine with one twin or the other. And this particular example indicates the general pattern: where the conflict account sees moral conflicts, the disjunctive account sees only disjunctive obligations.

In fact, the ideas underlying the conflict account described here are familiar, going back to Bas van Fraassen’s [1973]. The present formulation of the disjunctive account is new, however, and deserves further discussion. As we can see from its application to the twin example, there are actually two separate elements to the disjunctive account set out here. The first is the idea that, even when both of two prima facie oughts are triggered and undefeated, it is not necessary to accept either of the corresponding all things considered oughts, if they conflict—as in this case, where neither of the all things considered claims that I ought to dine with twins 1 or 2 is supported. The second element is the idea that a disjunction of the conflicting claims should be accepted as an all things considered ought—in this case, that I ought to dine with one or another of the two twins.

A view that seems to contain the first of these two elements without the second was proposed by Earl Conee, who agrees that there are cases in which “competing moral considerations have exactly the same force,” but writes that

there is no need to count each of these alternatives as morally required. We have the familiar option of holding that when moral factors are equal, each act is permitted and none is absolutely obligatory. [Conee, 1982, pp. 243–244]

What Conee suggests here is that, in a sense, the two counterbalanced moral claims cancel each other out, so that neither of the conflicting acts is obligatory, in accord with the first element of the disjunctive account; but although each of these acts is permitted, there appears to be no hint of the second element of the disjunctive account, according to which one of the two conflicting acts must be performed. A similar approach, containing the first but not the second element of the disjunctive account, was advanced by Foot, who considers a situation in which there are undefeated reasons for feeling that one ought to perform each of two incompatible actions, *a* and *b*; but rather than supposing that “both judgments have to be affirmed,” as the conflict account would have it, she is instead reluctant to draw either conclusion:

What we must ask, therefore, is whether in cases of irresolvable moral conflict we have to back both the judgment in favor of *a* and the judgment in favor of *b*, although doing *b* involves not doing *a*. Is it not possible that we should rather declare that the two are incommensurable, so that we have nothing to say about the overall merits of *a* and *b*. . . . [Foot, 1983, p. 267]

Again, Foot’s idea seems to be that we should refrain both from the judgment that one ought to perform the action *a* and from the judgment that one ought to perform the action *b*, but there is no suggestion that we should endorse the disjunctive judgment that one ought to perform either *a* or *b*.

As far as I know, this second element of the disjunctive account was first explicitly advanced by Donagan, in the course of commenting on an example involving conflicting but symmetrical *prima facie* oughts, like our dining example, but somewhat more dramatic:

Where the lives of identical twins are in jeopardy and I can save one but only one, every serious rationalist moral system lays down that, whatever I do, I must save one of them. . . . Certainly there is no moral conflict: from the fact that I have a duty to save either *a* or *b*, it does not follow that I have a duty to save *a* and a duty to save *b*. Can it be seriously held that a fireman, who has rescued as many as he possibly could of a group trapped in a burning building, should blame himself for the deaths of those left behind . . . ? [Donagan, 1984, pp. 286–287]

Still, although this passage does seem to contain a clear expression of the disjunctive idea, it is advanced only in terms of a particular example, from a

very different, rationalist perspective; and no precise account is provided of the way in which the output duties are supposed to be derived from the input rules supplied by Donagan's rationalist moral system.

It was not until Brink's [1994] that the disjunctive account received a full-scale defense from the present perspective, where the all things considered oughts are thought of as derived from an underlying set of prima facie oughts, without any particular rationalist constraints on the nature of these prima facie oughts. As in the present paper, Brink supposes that these all things considered oughts are to be generated from the undefeated prima facie oughts, but he rejects the view that each undefeated prima facie oughts should generate a corresponding all thing considered ought—a view that can be seen as a rudimentary version of our conflict account. Instead, he endorses an outcome, at least, that coincides with that provided by the disjunctive account as defined here.

Ordinarily, an undefeated prima facie obligation does constitute an all-things-considered obligation. But not always. Where there is an undefeated competitor, we can conclude that neither obligation is an all-things-considered obligation. This may seem to leave the agent confronting an insoluble conflict with no all-things-considered obligations, and this may seem puzzling to some. But the agent does face an all-things-considered obligation; it is to perform one or the other of the prima facie obligations. [Brink, 1994, p. 238]

Furthermore, unlike Donagan, Brink actually goes on to propose a procedure for specifying this desired outcome, deriving the all things considered oughts from the prima facie oughts. However, the procedure proposed by Brink is different from that set out here, and yields a result that fails to agree, I believe, both with that of the disjunctive account as defined here and with Brink's own desired outcome. In the present framework, as we recall, the disjunctive account is defined by appeal to the set of binding prima facie oughts. This is, of course, the same set that figures in the definition of the conflict account; the sole difference between the two accounts is that, rather than supporting a conclusion whenever it is entailed by some maximal consistent subset of these binding oughts, as in the conflict account, the disjunctive account requires that a supported conclusion should be entailed by every maximal consistent subset. On Brink's approach, by contrast, the set of binding oughts, those that are triggered but not defeated, is bypassed in favor of a different set of prima facie oughts: those that are triggered and, in addition to not being defeated themselves, also defeat all others with which they conflict.

An all-things-considered obligation represents what one ought to do in light of *all* morally relevant factors, including alternatives. If so, then only prima facie obligations that are undefeated *and* defeat all competitors are all-things-considered obligations. In other words, to be an all-things considered obligation, a prima facie obligation must be *overriding* and not simply not overridden. [Brink, 1994, p. 240]

Adapting Brink's terminology, and formalizing the notion of an overriding ought in the present setting, we can suppose that the oughts from the background context $\langle \mathcal{I}, \leq \rangle$ that are both triggered under the condition A , and that also defeat any other conflicting triggered oughts, are gathered together into the set $Overriding_{\langle \mathcal{I}, \leq \rangle}(A)$, defined as follows:

$$Overriding_{\langle \mathcal{I}, \leq \rangle}(A) = \{i \in \mathcal{I} : (1) i \in Triggered_{\mathcal{I}}(A), \\ (2) j < i \text{ for each } j \in Triggered_{\mathcal{I}}(A) \\ \text{such that } Consequent[i] \text{ and} \\ Consequent[j] \text{ are inconsistent}\}.$$

This definition should be compared to our previous definition of the binding prima facie oughts. As before, part (1) of the present definition tells us simply that each overriding prima facie ought must be triggered. But while part (2) of the previous definition tells us simply that a binding prima facie ought must be at least as important as any other ought with which it conflicts, what the present part (2) tells us is that an overriding prima facie ought must actually be more important than any conflicting ought.

It is easy to verify that, as Brink notes, the various prima facie oughts belonging to the set $Overriding_{\langle \mathcal{I}, \leq \rangle}(A)$ will be jointly consistent, and so the idea is that each of the overriding prima facie oughts from this set will give rise to an all things considered ought. More exactly, taking logical entailments into consideration, Brink's approach would suggest that an all things considered ought of the form $\bigcirc(B/A)$ should be taken as a consequence of the background context $\langle \mathcal{I}, \leq \rangle$ whenever B is logically entailed by the set $Consequent[Overriding_{\langle \mathcal{I}, \leq \rangle}(A)]$, containing the consequents of all the overriding prima facie oughts that are relevant under the condition A .⁶

In order to see the difficulty with this way of formulating the disjunctive standpoint, we need only to return to our twin example. Although both of the prima facie oughts i_3 and i_4 are triggered under the condition $A_1 \wedge A_2$, in which I have arranged to have dinner with both twins, neither of these two conflicting oughts actually defeats the other, since they are equally important. The set $Overriding_{\langle \mathcal{I}, \leq \rangle}(A_1 \wedge A_2)$, containing the overriding oughts, is therefore empty, as is the set $Consequent[Overriding_{\langle \mathcal{I}, \leq \rangle}(A_1 \wedge A_2)]$ containing the consequents of these overriding oughts, of course. As a result, it seems to follow from Brink's approach that an all things considered ought of the form $\bigcirc(B/A_1 \wedge A_2)$ should be supported in the twin example only when the statement B is a logical truth, so that, in particular, the disjunctive ought $\bigcirc(D_1 \vee D_2/A_1 \wedge A_2)$ —telling me that I ought to have dinner with one of the two twins—would not be supported.

Brink's own definitional procedure, then, appears to capture only the first, not the second, element of the disjunctive account that he advocates—successfully avoiding a conflict among all things considered oughts, but

failing to generate the appropriate disjunctive oughts. I conclude that the definition of the disjunctive account set out in the present paper, according to which a conclusion is supported whenever it is entailed by each maximal consistent subset of the binding oughts, provides a better match with the desired outcome, and will rely on this treatment in what follows.

2.3 Properties of the logics

Although not a central concern of this paper, it will be useful for the sake of perspective to note some formal properties of the two logics defined here for deriving all things considered oughts from a background context of prima facie oughts. As a preliminary step, we introduce the symbols \vdash_C and \vdash_D to represent the consequence relations defined by these two logics: the statements $\langle \mathcal{I}, \leq \rangle \vdash_C \bigcirc (B/A)$ and $\langle \mathcal{I}, \leq \rangle \vdash_D \bigcirc (B/A)$ will be taken to mean that the all things considered ought $\bigcirc(B/A)$ follows as a consequence of the background context $\langle \mathcal{I}, \leq \rangle$ according to the conflict or disjunctive accounts, respectively. We use the unadorned symbol \vdash —as in the statement $\langle \mathcal{I}, \leq \rangle \vdash \bigcirc (B/A)$ —when we wish to speak of both the conflict and disjunctive accounts indiscriminately.

The first thing to note about the two logics set out here is that neither allows for *strengthening of the antecedent*. Although it may be reasonable to conclude from the background context $\langle \mathcal{I}, \leq \rangle$ that a formula B ought to hold under a set of circumstances characterized only through the formula A , it need not follow that B ought to hold when the circumstances are characterized more fully through the formula $A \wedge C$ —or, put formally, from the fact that $\langle \mathcal{I}, \leq \rangle \vdash \bigcirc (B/A)$, it need not follow that $\langle \mathcal{I}, \leq \rangle \vdash \bigcirc (B/A \wedge C)$. The point can be illustrated through our earlier example, in which I have promised to meet a friend for lunch but am also needed to rescue a child from drowning. As we recall, the information in this example is represented through the context $\langle \mathcal{I}, \leq \rangle$, with $\mathcal{I} = \{i_1, i_2\}$ —where $i_1 = !(M/P)$ is the prima facie ought to meet my friend given my promise, and $i_2 = !(R/N)$ is the prima facie ought to rescue the child given the need—ordered so that $i_1 < i_2$; rescuing the child is strictly more important than meeting my friend. In this case, both of the two accounts developed here tell us that $\langle \mathcal{I}, \leq \rangle \vdash \bigcirc (M/P)$; if the situation is described only as one in which I have promised to meet my friend, it is reasonable to conclude that I ought to do so. But it does not follow from this that $\langle \mathcal{I}, \leq \rangle \vdash \bigcirc (M/P \wedge N)$. When the situation is described as one in which I have promised to meet my friend but am also needed to rescue the child, we no longer conclude that I ought to meet my friend; instead, both of the two accounts now tell us $\langle \mathcal{I}, \leq \rangle \vdash \bigcirc (R/P \wedge N)$, that I ought to rescue the child.

In failing to allow for strengthening of the antecedent, the accounts set out here agree with those logics of conditional obligation that are developed as a species of conditional logic, within the framework of possible worlds semantics.⁷ But there is another, deeper way in which the accounts developed here

differ even from these conditional logics—namely, in failing to satisfy the property of *consequence monotonicity*. In classical logic, as well as most standard philosophical logics, including conditional logic, the set of derivable conclusions grows monotonically along with the information contained in a set of premises: increasing the information contained in a premise set will never force the abandonment of a previously supported conclusion—or, put more formally, if the premise set Γ contains all the information found in the premise set Γ' , we know that a formula A will be a consequence of Γ whenever A is a consequence of Γ' . If we follow the natural route of taking background contexts as analogous to premise sets, however, this property of consequence monotonicity fails in the case of the present logics. Here, a context $\langle \mathcal{I}, \leq \rangle$ can be thought of as containing at least as much information as a context $\langle \mathcal{I}', \leq' \rangle$ whenever all the prima facie oughts from \mathcal{I}' are included among those in \mathcal{I} and all the ordering information from \leq' is included among that in \leq . Yet it is possible, in the present case, for a context $\langle \mathcal{I}', \leq' \rangle$ to support as a conclusion an all things considered ought that is not supported by a context $\langle \mathcal{I}, \leq \rangle$, even though $\langle \mathcal{I}, \leq \rangle$ contains at least as much information as $\langle \mathcal{I}', \leq' \rangle$.

Again, this point can be illustrated through the same example, involving the clash between meeting a friend for lunch and rescuing a child. This time, however, let us suppose that $\langle \mathcal{I}', \leq' \rangle$ is my initial background context, where the initial set $\mathcal{I}' = \{i_1\}$ of prima facie oughts contains only one of the two imperatives from the original example, concerning my promise to meet my friend, and where the initial ordering \leq' tells us only, trivially, that this prima facie ought is at least as important as itself. In this case, we have $\langle \mathcal{I}', \leq' \rangle \vdash \sim \bigcirc (M/P \wedge N)$; both accounts tell us that I ought to meet my friend, even in light of the need to rescue the child, since, of course, there is no prima facie ought in my background context that is triggered by this need. Now let us consider again the original background context $\langle \mathcal{I}, \leq \rangle$, with $\mathcal{I} = \{i_1, i_2\}$ containing both imperatives, and \leq defined so that $i_1 < i_2$, telling us again that rescuing a child is more important than meeting a friend. Here, it is clear that the context $\langle \mathcal{I}, \leq \rangle$ contains all the information—all the prima facie oughts and ordering constraints—present in the context $\langle \mathcal{I}', \leq' \rangle$. Nevertheless, we no longer have $\langle \mathcal{I}, \leq \rangle \vdash \sim \bigcirc (M/P \wedge N)$ since the context $\langle \mathcal{I}, \leq \rangle$ does now contain a prima facie ought that is triggered by the need to rescue the child, and that, in fact, defeats the imperative telling me to meet my friend.

Because the addition of new information to a background context can lead, in the accounts set out here, to the abandonment of previously supported conclusions, these two accounts fail to satisfy the property of consequence monotonicity. It follows that these accounts cannot be articulated in any simple way within the modal, or intensional, framework that is so often appealed to as a formal foundation for deontic logic, since theories developed within this framework tend to support consequence monotonicity. Instead, a formal development of the ideas sketched here would most naturally involve techniques from the field of nonmonotonic logic.⁸

Still, even though a full development of the conflict and disjunctive accounts would have to appeal to new and somewhat nonstandard logical techniques, it is interesting to note that both of these two accounts are *conservative extensions* of standard deontic logic in the sense that, roughly speaking, when applied to a background context of prima facie oughts that is conflict free, these two accounts can be seen to agree both with standard deontic logic and with each other. The two accounts differ, from standard deontic logic and from each other, only when applied to a background context containing conflicting prima facie oughts; and in this case, the conflict account defined here itself can be seen as a conservative extension of the approach originally set out by van Fraassen [1973]. Both of these results—that the conflict and disjunctive accounts are both conservative extensions of standard deontic logic, and that the conflict account is itself a conservative extension of van Fraassen’s approach—are formulated precisely and established in an appendix to this paper.

Although the definitions of consequence for the conflict and disjunctive accounts presented thus far specify the notion only for atomic ought statements of the form $\bigcirc(B/A)$, the idea can be extended to Boolean combinations of these statements in the obvious way, by stipulating that

$$\begin{aligned} \langle \mathcal{I}, \leq \rangle \sim \phi \wedge \psi & \text{ if and only if } \langle \mathcal{I}, \leq \rangle \sim \phi \text{ and } \langle \mathcal{I}, \leq \rangle \sim \psi, \\ \langle \mathcal{I}, \leq \rangle \sim \neg\phi & \text{ if and only if it is not the case that } \langle \mathcal{I}, \leq \rangle \sim \phi, \end{aligned}$$

for combinations ϕ and ψ of atomic oughts. It is then easy to verify that both of the resulting logics defined here are *noncontradictory*, in the sense that neither will allow a consequence statement of the form $\langle \mathcal{I}, \leq \rangle \sim \phi \wedge \neg\phi$. No background context, no matter what its nature, will ever support a contradiction.

It is particularly important, in the case of the conflict account, to keep carefully in mind the distinction between moral conflicts, of the sort that are allowed by this account, and logical contradictions, which are forbidden. According to the conflict account, it is possible for a background context $\langle \mathcal{I}, \leq \rangle$ to support two oughts of the form $\bigcirc(B/A)$ and $\bigcirc(\neg B/A)$, telling us both that we ought to have B under the condition A , and also that we ought to have $\neg B$ under the condition A ; this is simply an all things considered moral conflict, of the sort that the conflict account is designed to allow. But since the account is noncontradictory, it can never support two statements of the form $\bigcirc(B/A)$ and $\neg \bigcirc(B/A)$, telling us both that we ought to have B under the condition A , and also that it is not the case that we ought to have B under the condition A ; that would be a contradiction.

Continuing our survey of logical properties, we can see also that both the conflict and disjunctive accounts defined here are characterized by a notion of consequence for all things considered oughts that is itself *closed under logical consequence*, in the sense that, whenever we know both that

$\langle \mathcal{I}, \leq \rangle \sim \circ (B/A)$ and that $B \vdash C$ —that is, C is an ordinary logical consequence of B —we must also have $\langle \mathcal{I}, \leq \rangle \sim \circ (C/A)$. The intuitive force of this property can be seen through consideration of a case in which $\langle \mathcal{I}, \leq \rangle \sim \circ (B/A)$ and $D \vdash \neg B$, where closure under logical consequence would lead us to conclude $\langle \mathcal{I}, \leq \rangle \sim \circ (\neg D/A)$. Here, the idea is that anything that entails the negation of what ought to be the case had better not be the case itself: since it ought to be the case that B , and D logically entails $\neg B$, it had better be the case that $\neg D$.⁹

Finally, it is easy to see that both the conflict and disjunctive accounts support the principle that *ought implies can*, sometimes known as the *volunteerist* principle, according to which only consistent formulas can be enjoined, regardless of the background set of prima facie oughts; put formally, what the principle tells us is that, whenever $\langle \mathcal{I}, \leq \rangle \sim \circ (B/A)$, it follows that the enjoined formula B must itself be consistent. In fact, however, these two accounts block the support of inconsistent all things considered oughts in interestingly different ways, reflected in their different treatments of a principle characterized by Bernard Williams [1965] as the rule of *agglomeration*, according to which any conjunctions formed from enjoined formulas must be enjoined as well.¹⁰ Formally, and in full generality, the rule of agglomeration allows us to conclude that $\langle \mathcal{I}, \leq \rangle \sim \circ (B \wedge C/A)$ whenever we know that $\langle \mathcal{I}, \leq \rangle \sim \circ (B/A)$ and $\langle \mathcal{I}, \leq \rangle \sim \circ (C/A)$, but it is perhaps more easily recognizable in its unconditional form, according to which it allows us to conclude $\langle \mathcal{I}, \leq \rangle \sim \circ (B \wedge C)$ from $\langle \mathcal{I}, \leq \rangle \sim \circ (B)$ and $\langle \mathcal{I}, \leq \rangle \sim \circ (C)$.

In order to illustrate the different ways in which the conflict and disjunctive accounts block the support of inconsistent oughts, and their different treatments of agglomeration, let us now consider an extreme case. Suppose the background context is $\langle \mathcal{I}, \leq \rangle$, where $\mathcal{I} = \{!(A), !(A)\}$ tells us that both A and $\neg A$, considered individually, ought prima facie to be the case, and where \leq ranks neither of these two oughts as more important than the other. Since both of these prima facie oughts are binding, the conflict account must promote both to all things considered oughts—that is, we must have both $\langle \mathcal{I}, \leq \rangle \sim_C \circ (A)$ and $\langle \mathcal{I}, \leq \rangle \sim_C \circ (\neg A)$. It is therefore plain that the consequence relation associated with the conflict account cannot allow the rule of agglomeration, for agglomeration would lead us to conclude at once, in this case, that $\langle \mathcal{I}, \leq \rangle \sim_C \circ (A \wedge \neg A)$, which violates the principle that ought implies can, of course, since the formula $A \wedge \neg A$ is inconsistent. As in the approach originally advocated by Williams, the conflict account allows for the support of a collection of all things considered oughts that, taken jointly, are inconsistent, but it blocks the support of individual inconsistent oughts by refusing to allow these jointly inconsistent oughts to be agglomerated into one.

The disjunctive account, by contrast, does allow the rule of agglomeration: in general, we can conclude that $\langle \mathcal{I}, \leq \rangle \sim_D \circ (B \wedge C/A)$ whenever we know that $\langle \mathcal{I}, \leq \rangle \sim_D \circ (B/A)$ and $\langle \mathcal{I}, \leq \rangle \sim_D \circ (C/A)$. In the case of our extreme

example, we would be able to conclude from agglomeration that $\langle \mathcal{I}, \leq \rangle \sim_D \circ (A \wedge \neg A)$ if we could establish both $\langle \mathcal{I}, \leq \rangle \sim_D \circ (A)$ and $\langle \mathcal{I}, \leq \rangle \sim_D \circ (\neg A)$. And so it follows at once, since the disjunctive account also satisfies the principle that ought implies can, that we should not be able to establish both $\langle \mathcal{I}, \leq \rangle \sim_D \circ (A)$ and $\langle \mathcal{I}, \leq \rangle \sim_D \circ (\neg A)$ —and indeed we can establish neither, since neither of these two conclusions is supported by each maximal consistent subset of the binding oughts. Unlike the conflict account, then, the disjunctive account does allow various supported all things considered oughts to be agglomerated into one, but in this case there is no risk that such agglomeration would lead to an individually inconsistent ought since the entire collection of supported oughts is itself jointly consistent.

3 Evaluating the approaches

We now have before us two logics, reflecting two different approaches to the problem of deriving all things considered oughts from a background context of prima facie oughts: the conflict account, according to which the all things considered oughts are those entailed by some maximal consistent subset of the binding prima facie oughts, and the disjunctive account, according to which an all things considered ought must be entailed by each maximal consistent subset of the binding prima facie oughts.

Although these two logics lead to different verdicts on the issue of all things considered moral conflict—one denying and the other affirming the possibility—they are developed within a common framework, using common ideas; and as we have seen, the resulting logics share a number of desirable properties. Both are conservative extensions of standard deontic logic, agreeing with this theory, and with each other, when applied to a consistent background context of prima facie oughts. Both are noncontradictory, never supporting the negation of a supported all things considered ought, regardless of the background context. Both support a set of all things considered oughts that is closed under logical consequence, enjoining all the logical consequences of the enjoined propositions. Both maintain the principle that ought implies can, according to which only consistent propositions can be enjoined. And both are sensitive enough to register the failure of strengthening in the antecedent, and also the failure of consequence monotonicity, that characterizes our normative reasoning.

Against this background, we can now turn to the main question at issue in this paper: setting aside arguments based on the adoption of some particular moral theory or another, are there any reasons for favoring one of these two approaches over the other, and in particular, for rejecting the conflict approach? We consider three kinds of objections to the possibility of all things considered moral conflicts, based respectively on logical arguments, conceptual considerations concerning the nature of all things considered oughts, and an analogy between moral and physical forces.

3.1 Logical objections: agglomeration

There are three standard logical objections to any theory, such as the conflict account, that allows for the possibility of conflicting all things considered oughts. Although these three objections are admirably presented, along with references to the literature, by Christopher Gowans [1987a], we will concentrate here on the more recent and extensive discussion found in Brink's [1994]. In this paper, each of these objections is presented as a "paradox"—a contradiction generated from the assumption of conflicting all things considered oughts, taken together with certain principles from standard deontic logic that are viewed as crucial to our normative reasoning. Since, as we have seen, the conflict account defined here is itself noncontradictory, any derived contradictions must depend in an essential way on these auxiliary principles, and so our assessment of these principles will likewise determine the force of the objections.

Brink's first paradox relies on two deontic principles, both discussed in the previous section. The first is the principle that ought implies can, which is already supported by the conflict account; the second is the rule of agglomeration, which is not. It is of course easy, as we have just seen, to arrive at a contradiction when the conflict account is supplemented with agglomeration. The conflict account allows certain background contexts to support statements of the form $\bigcirc(A)$ and $\bigcirc(\neg A)$ —that is, we might have both $\langle \mathcal{I}, \leq \rangle \sim_C \bigcirc(A)$ and $\langle \mathcal{I}, \leq \rangle \sim_C \bigcirc(\neg A)$. Agglomeration would then lead us to the conclusion $\langle \mathcal{I}, \leq \rangle \sim_C \bigcirc(A \wedge \neg A)$, enjoining the formula $A \wedge \neg A$. The principle that ought implies can tells us that any enjoined formula must be consistent. But the formula $A \wedge \neg A$ is not consistent.

In fact, Brink himself does not take this first paradox very seriously, since he tends to dismiss the principles through which it is generated, and agglomeration in particular.

Consider agglomeration. Where there is no conflict between A and B , it seems harmless to recognize an obligation to do both as well as obligations to do each. But it also seems unnecessary; an obligation to do each seems adequate to explain the moral situation. [Brink, 1994, p. 229]

I believe, however, that the matter is more complicated, and that we do need to allow for some degree of agglomeration in order to account for certain aspects of our moral reasoning, even in an account that tolerates conflicts.

The point is best illustrated with an example. Imagine that an agent's background moral theory contains only the two *prima facie* oughts "You ought either to fight in the army or perform alternative service" and "You ought not to fight in the army"—the first issuing from some legal authority, perhaps, the second from religion or conscience. More exactly, if we let F represent the proposition that the agent should fight in the army and S the

proposition that the agent should perform alternative service, so that the two prima facie oughts are

$$i_5 = !(F \vee S),$$

$$i_6 = !(\neg F),$$

we can take the context as $\langle \mathcal{I}, \leq \rangle$, with $\mathcal{I} = \{i_5, i_6\}$ and \leq arbitrary. Since these two prima facie oughts are consistent, they should, on any account, be promoted to the respective all things considered oughts $\bigcirc(F \vee S)$ and $\bigcirc(\neg F)$, telling us that the agent ought, all things considered, to fight in the army or perform alternative service, but also that he ought not to fight in the army. Is this a complete description of the moral situation? I believe not. I believe that, once we reach the conclusions that the agent ought either to fight in the army or perform alternative service, but also that he ought not to fight in the army, we are then committed to the further conclusion that the agent ought to perform alternative service—that is, to the further conclusion $\bigcirc(S)$.

Of course, one might argue that it is not necessary for the agent to derive this further conclusion explicitly in order to perform the right action, and perhaps this is Brink's point. As long as the agent satisfies the explicitly presented all things considered oughts $\bigcirc(F \vee S)$ and $\bigcirc(\neg F)$, bringing about a situation in which he either fights in the army or performs alternative service, but also one in which he does not fight in the army, the agent will satisfy the implicitly derivable ought $\bigcirc(S)$ as well, bringing about a situation in which he performs alternative service. Since this ought will be satisfied in any case, why, then, is it necessary for it to be explicitly derived? To argue in this way, however, would be to limit the scope of deontic logic to a narrowly action-guiding enterprise, rather than one that is supposed to be more fully descriptive of the moral situation. If the formalism is to serve simply as a guide to action, it may be sufficient for it to enjoin a set of formulas which, as long as these formulas are satisfied, will lead to the achievement of a proper state of affairs. If the aim is descriptive, on the other hand, it is natural to expect a deontic logic to provide a more complete characterization of the moral situation; and in the case at hand, it seems clear that such a characterization should include the information that one of the things the agent ought to do, all things considered, is perform alternative service.

The advantage of the rule of agglomeration is that, together with closure under logical consequence, it allows us to reach conclusions like this, which involve logical interactions among enjoined formulas. In the present case, if we were able to combine the enjoined formulas from the two all things considered oughts $\bigcirc(F \vee S)$ and $\bigcirc(\neg F)$ into the agglomerated result $\bigcirc((F \vee S) \wedge \neg F)$, we could then immediately derive $\bigcirc(S)$ as a conclusion, given closure under consequence, since S is a logical consequence of $(F \vee S) \wedge \neg F$. Of course, as we have seen, full agglomeration cannot coherently be accommodated within the conflict account, unless we are willing to

abandon the principle that ought implies can. But without some measure of agglomeration, at least, it is difficult to see how we could reach certain desirable results, such as the present conclusion that the agent ought to perform alternative service.¹¹ Apparently, what is needed is some degree of agglomeration, but not too much; and the problem of formulating a principle allowing for exactly the right amount of agglomeration raises delicate issues that have generally been ignored in the literature, which seems to contain only arguments favoring either wholesale acceptance or wholesale rejection.

The conflict account provides an example of such a hedged agglomeration principle, falling between wholesale acceptance and wholesale rejection. As we have seen, this account avoids generating unfortunate agglomerates of the form $\bigcirc(A \wedge \neg A)$, even when both $\bigcirc(A)$ and $\bigcirc(\neg A)$ are supported. But it is easy to see also, in the present example, that the account does support the agglomerate $\bigcirc((F \vee S) \wedge \neg F)$, and so the desired result $\bigcirc(S)$. In this case, $\text{Binding}_{\langle \mathcal{I}, \leq \rangle}(\top) = \{i_5, i_6\}$ is the set of prima facie oughts that are binding under the circumstances, so that $\text{Consequent}[\text{Binding}_{\langle \mathcal{I}, \leq \rangle}(\top)] = \{F \vee S, \neg F\}$ is the set of their consequents; but since this set is consistent, it is its own unique maximal consistent subset, and entails both the formulas $((F \vee S) \wedge \neg F)$ and S .

The style of agglomeration sanctioned by the conflict account can be characterized as *consistent consequent agglomeration*, and defined as follows. Suppose we know that a number of oughts of the form $\bigcirc(B_i/A)$ are supported by some background context $\langle \mathcal{I}, \leq \rangle$ —that is, we have $\langle \mathcal{I}, \leq \rangle \sim \bigcirc(B_i/A)$ for i from 1 through n . Then we can conclude that the agglomerate $\bigcirc(B_1 \wedge B_2 \wedge \dots \wedge B_n/A)$ is likewise supported—that is, $\langle \mathcal{I}, \leq \rangle \sim \bigcirc(B_1 \wedge B_2 \wedge \dots \wedge B_n/A)$ —as long as we know that the set $\{B_1, B_2, \dots, B_n\}$ is both consistent and also a subset of the set $\text{Consequent}[\text{Binding}_{\langle \mathcal{I}, \leq \rangle}(A)]$, containing the consequents of the binding prima facie oughts. Although the statement of this rule is somewhat complex, the underlying idea is simple. What it tells us is that we can agglomerate a set of oughts into a conjunctive ought under two conditions: first, the set of formulas enjoined by these various oughts must be consistent; second, each of these enjoined formulas must be the consequent of some binding prima facie ought.

This second condition—that each enjoined formula should be the consequent of some binding prima facie ought—may seem peculiar, or at least excessively syntactic. Dropping this restriction would result in another hedged agglomeration principle, stronger than consistent consequent agglomeration but still weaker than full agglomeration, that could be characterized simply as *consistent agglomeration*. Formally, such a rule would allow us to conclude that $\langle \mathcal{I}, \leq \rangle \sim \bigcirc(B \wedge C/A)$ whenever we know that $\langle \mathcal{I}, \leq \rangle \sim \bigcirc(B/A)$ and $\langle \mathcal{I}, \leq \rangle \sim \bigcirc(C/A)$, subject only to the constraint that the set $\{B, C\}$ should be consistent.¹² Again, this rule—which tells us that we are free to agglomerate individual oughts as long as doing so does not actually lead to an inconsistent enjoined formula—is weak enough to avoid generating agglomerates of the form $\bigcirc(A \wedge \neg A)$ even when both of

the individual oughts $\bigcirc(A)$ and $\bigcirc(\neg A)$ are supported; and it is strong enough, in the present example, to tell us that we are committed to the agglomerate $\bigcirc((F \vee S) \wedge \neg F)$, and so to $\bigcirc(S)$, once we have accepted the individual oughts $\bigcirc(F \vee S)$ and $\bigcirc(\neg F)$.

It turns out, however, that this stronger rule of consistent agglomeration is not supported by the conflict account. We can see this, and also understand the point of the syntactic restrictions that distinguish the rule of consistent consequent agglomeration, by considering yet another example, based on a familiar problem. Suppose that both you and your spouse have a prima facie obligation to visit your own parents during the holiday season, but that both of you likewise have a prima facie obligation to visit your spouse's parents. And suppose that, because they live in separate parts of the country, it is simply not possible to visit both sets of parents; but of course, whichever set of parents you decide on, they should first be notified of your plans to visit. Let us take N_1 as the proposition that you will notify your own parents of your plans to visit during the holidays, V_1 as the proposition that you will actually visit your own parents, and N_2 and V_2 as the respective propositions that you will notify and visit your spouse's parents. Your prima facie obligations can then be represented through the two imperatives

$$\begin{aligned} i_7 &= !(N_1 \wedge V_1), \\ i_8 &= !(N_2 \wedge V_2), \end{aligned}$$

telling us that you should notify and then visit your own parents, but also that you should notify and then visit your spouse's parents; the background context is therefore $\langle \mathcal{I}, \leq \rangle$, where $\mathcal{I} = \{i_7, i_8\}$ and we can assume that \leq ranks i_7 and i_8 as either equal or incomparable in importance.

In this case, both of the prima facie oughts are binding, of course—that is, we have $\text{Binding}_{\langle \mathcal{I}, \leq \rangle}(\top) = \{i_7, i_8\}$, so that $\text{Consequent}[\text{Binding}_{\langle \mathcal{I}, \leq \rangle}(\top)] = \{N_1 \wedge V_1, N_2 \wedge V_2\}$ is the set of their consequents. Since it is impossible to visit both your own parents and your spouse's parents, it is natural to view the statements V_1 and V_2 as inconsistent. The entire set of consequents is therefore inconsistent as well, with $\mathcal{F}_1 = \{N_1 \wedge V_1\}$ and $\mathcal{F}_2 = \{N_2 \wedge V_2\}$ as its maximal consistent subsets. Since \mathcal{F}_1 entails V_1 and \mathcal{F}_2 entails V_2 , the conflict account tells us that you ought to visit your own parents, and also that you ought to visit your spouse's parents—that is, we have both $\langle \mathcal{I}, \leq \rangle \sim_C \bigcirc(V_1)$ and $\langle \mathcal{I}, \leq \rangle \sim_C \bigcirc(V_2)$. But since the set $\{V_1, V_2\}$ is inconsistent, neither of the two rules under consideration, consistent agglomeration or consistent consequent agglomeration, entails the conclusion $\langle \mathcal{I}, \leq \rangle \sim_C \bigcirc(V_1 \wedge V_2)$. This is, of course, as it should be: the conflict account does not support the conclusion that you ought to visit both sets of parents, since doing so is impossible.

The conflict account also tells us, however, both that you ought to notify your own parents that you will visit and that you ought to notify your

spouse's parents—again, since \mathcal{F}_1 entails N_1 and \mathcal{F}_2 entails N_2 , we have both $\langle \mathcal{I}, \leq \rangle \vdash_C \bigcirc(N_1)$ and $\langle \mathcal{I}, \leq \rangle \vdash_C \bigcirc(N_2)$. And here, it is natural to view the set $\{N_1, N_2\}$ as consistent, since, even though you will not actually be able to visit both sets of parents, it is certainly possible for you to notify both set of parents that you will be visiting. In this case, the rule of consistent agglomeration would therefore suggest the conclusion $\langle \mathcal{I}, \leq \rangle \vdash_C \bigcirc(N_1 \wedge N_2)$, which is incorrect: the conflict account does not tell us that you should notify both sets of parents of an impending visit. It is exactly this kind of odd result that the rule of consistent consequent agglomeration manages to avoid, through its syntactic restrictions. Even though the set $\{N_1, N_2\}$ is consistent, it is not a consistent subset of the set $\text{Consequent}[\text{Binding}_{\langle \mathcal{I}, \leq \rangle}(\top)] = \{N_1 \wedge V_1, N_2 \wedge V_2\}$, containing the consequents of the binding oughts; and of course, no consistent subset of this latter set entails the formula $N_1 \wedge N_2$.

Within the context of the conflict account, the rule of consistent consequent agglomeration seems to provide all the agglomeration we need and no more than we want. Although I do not believe the matter is entirely settled by the arguments presented here, the rule seems to hold some promise, at least, as a correct, appropriately hedged principle of agglomeration.

3.2 Other logical objections

Let us now turn our attention from considerations concerning the proper form of an appropriate agglomeration principle to the remaining two logical arguments against a theory that allows all things considered moral conflicts—Brink's remaining two paradoxes, which he does take seriously, unlike the first, since he endorses the principles involved.

The second logical objection is again based on two principles. The first is the principle of closure under logical consequence, which Brink refers to as the *obligation execution* principle, because it obliges us not to do anything that would prevent the execution of our obligations. This principle, which is supported by the conflict account, was presented in the previous section as the principle allowing us to conclude, whenever we know that $\langle \mathcal{I}, \leq \rangle \vdash \bigcirc(B/A)$ and $B \vdash C$ —that is, C is an ordinary logical consequence of B —that $\langle \mathcal{I}, \leq \rangle \vdash \bigcirc(C/A)$ should hold as well; but for the purpose of the present discussion, it will be sufficient to concentrate only on an unconditional version of the principle, allowing us to conclude, from $\langle \mathcal{I}, \leq \rangle \vdash \bigcirc(B)$ and $B \vdash C$, that $\langle \mathcal{I}, \leq \rangle \vdash \bigcirc(C)$. The second principle, which Brink describes as the *weak obligation* principle, tells us that, whenever a formula B is enjoined as an all things considered ought, it is not the case that $\neg B$ can be enjoined as well—that is, that $\bigcirc(B)$ implies $\neg \bigcirc(\neg B)$. In the present setting, this principle can be captured as the assumption that

$$\langle \mathcal{I}, \leq \rangle \vdash \bigcirc(B) \supset \neg \bigcirc(\neg B)$$

should hold quite generally, for any background context $\langle \mathcal{I}, \leq \rangle$ and formula B whatsoever.

It is easy to see the troubles that result when the conflict account, which does not support the weak obligation principle, is then supplemented with this principle. Suppose that, as the conflict account allows, some context enjoins both of two inconsistent formulas as all things considered oughts—that is, suppose that (1) $\langle \mathcal{I}, \leq \rangle \sim_C \bigcirc (G)$ and (2) $\langle \mathcal{I}, \leq \rangle \sim_C \bigcirc (H)$, where the formulas G and H are inconsistent. Since G and H are inconsistent, we can of course conclude (3) $H \vdash \neg G$ —that H logically implies the negation of G . From (2) and (3), the principle of closure under logical consequence, or obligation execution, allows us to conclude that (4) $\langle \mathcal{I}, \leq \rangle \sim_C \bigcirc (\neg G)$. From (1), however, the principle of weak obligation yields the conclusion (5) $\langle \mathcal{I}, \leq \rangle \sim_C \neg \bigcirc (\neg G)$.¹³ And then, together, (4) and (5) show that the conflict account, when it is supplemented with the weak obligation principle, leads to a contradiction—not just to the idea that conflicting formulas might both be enjoined as all things considered oughts, which is simply a moral conflict, but to the idea that the very same formula might be both enjoined and not enjoined, which is a logical contradiction.

The third of the standard logical objections to moral conflicts—Brink's third paradox—relies on the new notion of permissibility, where the formula $P(B)$ is taken to represent the idea that the proposition B is permissible. Typically in deontic logic, what is permissible is characterized in terms of what ought to be the case, through the definition

$$P(B) =_{df} \neg \bigcirc (\neg B),$$

telling us that a proposition is permissible if it is not the case that it ought not to be the case. Brink accepts this characterization of permissibility in terms of ought, not as a definition, actually, but rather as a principle relating the two ideas, which he refers to as the principle of *correlativity*. In addition, however, Brink also accepts the further and more substantive principle, which he describes as the *weak impermissibility* principle, according to which anything that ought to be the case must also be permissible—that is, according to which $\bigcirc(B)$ implies $P(B)$.¹⁴ This principle is best captured in the present setting through the general stipulation that

$$\langle \mathcal{I}, \leq \rangle \sim \bigcirc (B) \supset P(B)$$

holds for each context and formula.

Again, the weak impermissibility principle is not supported by the conflict account, and as before, it is easy to see the difficulties that result when the account is supplemented with this principle by considering a case in which some context enjoins conflicting formulas as all things considered oughts—that is, where (1) $\langle \mathcal{I}, \leq \rangle \sim_C \bigcirc (G)$ and (2) $\langle \mathcal{I}, \leq \rangle \sim_C \bigcirc (H)$, where G and H are

inconsistent. Since G and H are inconsistent, we again have (3) $H \vdash \neg G$. From (2) and (3), the principle of closure under consequence again tells us that (4) $\langle \mathcal{I}, \leq \rangle \sim_C \bigcirc(\neg G)$. But now, from (1), Brink's weak impermissibility principle tells us that (5) $\langle \mathcal{I}, \leq \rangle \sim_C \bigcirc P(G)$, which is equivalent to (6) $\langle \mathcal{I}, \leq \rangle \sim_C \neg \bigcirc(\neg G)$ by the definition of what is permissible in terms of what ought to be. And then (4) and (6) show that the conflict account, supplemented with the weak impermissibility principle, again leads to a contradiction.

What Brink's second and third logical objections show, then, is that the conflict account is rendered inconsistent if it is supplemented with either the weak obligation principle or the weak impermissibility principle, both of which he supports. But how convincing are these objections, taken as arguments against the conflict account? The first thing to note is that it is not really surprising that these principles should lead to inconsistency in the context of the conflict account—which allows for moral conflicts, of course—since the effect of each of these principles is simply to assert that there can be no moral conflicts. The weak obligation principle, for example, supplements the set of statements supported by the conflict account with arbitrary formulas of the form $\bigcirc(B) \supset \neg \bigcirc(\neg B)$, but any such formula is equivalent by elementary logic to a formula of the form $\neg \bigcirc(B) \vee \neg \bigcirc(\neg B)$, which is then equivalent to the formula $\neg(\bigcirc(B) \wedge \bigcirc(\neg B))$ —telling us explicitly that the conflicting formulas B and $\neg B$ cannot both be enjoined. And the weak impermissibility principle supplements the set of supported statements with formulas of the form $\bigcirc(B) \supset P(B)$, but each of these can then be transformed, by the definition of permissibility in terms of ought, into a weak obligation statement of the form $\bigcirc(B) \supset \neg \bigcirc(\neg B)$, which can itself be transformed into a statement explicitly denying the existence of moral conflicts, as we have just seen.

In some ways, then, since the principles of weak obligation and weak impermissibility are both equivalent to the denial of moral conflicts, it is hard to take them seriously as components of any real argument against the existence of such conflicts, as opposed simply to statements of the view that there are none; both principles are nothing but straightforward denials of the very point at issue. Still, restating a view in a particular way sometimes highlights advantages to it that may not have been apparent in its original formulation, even when the original formulation and the restatement are equivalent—and perhaps the reason for denying the possibility of moral conflicts is more evident when this position is restated as the weak obligation principle, or as weak impermissibility. It is therefore important to consider the arguments that Brink actually presents in favor of the weak obligation and weak impermissibility principles; but in both cases, the justification is brief, as if the truth of these principles should be almost immediately apparent. Concerning weak obligation, he says only:

And the weak obligation principle, as its name suggests, seems especially uncontroversial. If I'm obligated not to kill my neighbor, then surely it's not the case that I'm obligated to kill him. [Brink, 1994, p. 235]

And concerning weak impermissibility:

The new deontic principle [in the third paradox] is the weak impermissibility principle. But surely that must be true. If it's impermissible for me to torture my neighbor, then surely it's not the case that I'm obliged to torture him. [Brink, 1994, p. 236]

What Brink provides here, by way of justification, are simply a couple of instances of the abstract principles in question that are supposed to seem intuitively attractive—which is perfectly fair, of course: at some point in the justification of a set of fundamental logical principles, appeal to argument must necessarily be replaced by appeal to intuition. Still, it is difficult to establish a general principle beyond doubt by exhibiting a single instance, no matter how convincing. There is always the possibility that the truth of that particular instance could be explained through some other means, rather than by appeal to the general principle involved, or that a plausible counterexample to the principle might still be found; and in this case it seems that both avenues of criticism are available.

Here, the instances provided by Brink to illustrate the abstract principles are so alarming, and the consequents of these particular instances—that I am not obligated to kill my neighbor, or to torture my neighbor—are so palpably true that it is hard to see whether these consequents actually follow from the antecedents of the principles, or whether they are simply true on their own, lending a false credibility to the principles themselves. Moreover, there are coherent scenarios in which the two principles could at least arguably be classified as incorrect, as we can see by recalling the situation in which I have promised to have dinner tonight with each of two identical twins, but cannot dine with both, so that dining with either twin entails not dining with the other. In this situation, it can be argued that, since I promised to dine with twin 2, I ought to do so, and therefore—given closure under consequence, because dining with twin 2 entails not dining with twin 1—that I ought not to dine with twin 1. But since I promised to dine with twin 1 as well, I ought to dine with twin 1. The situation is therefore arguably one in which I ought to dine with twin 1 but also ought not to dine with twin 1, contrary to the principle of weak obligation. Furthermore, the definition of permissibility in terms of ought tells us that I am not permitted to dine with twin 1 if I ought not to dine with twin 1. And so the same situation—in which I ought to dine with twin 1 but also ought not to dine with twin 1—arguably illustrates a case in which I ought to dine with twin 1 but,

according to the definition, am not permitted to do so, contrary to the principle of weak impermissibility.

3.3 *Conceptual objections*

The present paper relies crucially on a conceptual distinction between two kinds of oughts. Oughts of the first kind, described here as *prima facie* oughts, specify only moral reasons for preferring one state of affairs or action over another; oughts of the second kind, described here as all things considered oughts, are supposed to reflect a more general judgment of desirability, which results when all of the various *prima facie* oughts are taken into consideration.

In broad outline, this distinction coincides with that drawn by any number of writers who also rely on some contrast between, let us say, *weak* and *strong* oughts, where the weak oughts provide reasons for reaching a moral judgment and the strong oughts register the judgments reached on the basis of those reasons. Although some of these writers—Roderick Chisholm [1964], Donald Davidson [1970]—use the vocabulary of *prima facie* oughts to characterize this distinction, others present the view using different terminology. For example, Searle [1980] refers to our weak oughts as “obligations,” and uses the term “ought” itself only for strong oughts; the idea is that agents are subject to various obligations, possibly conflicting, which are then combined to result in a specification of what, all things considered, they ought to do. In a similar fashion, Thomson [1990] describes our weak oughts in terms of the “obligations” and “commitments” we have to one another, or—echoing the language of Wesley Hohfeld—the “claims” we have against each other, and again reserves the term “ought” itself for the strong oughts that result when all of these weak oughts are taken into account. Other writers, such as Gowans [1987a], actually use the term “ought” to refer to our weak oughts, referring to our strong oughts as what we “must” do. And Foot [1983] abandons any attempt to link the distinction between weak and strong oughts to familiar terms from natural language, speaking of them only as “type 1” and “type 2” oughts respectively.

All of these writers, then, agree with the present treatment in accepting the contrast between weak and strong oughts, in one guise or another. Where they differ, however, is on the following crucial point. The present paper sets out two accounts of the all things considered, or strong, ought—one of which, the conflict account, allows for moral conflicts even among strong oughts. These various writers, on the other hand, once they have drawn the distinction between weak and strong oughts, either argue or, more often, simply assume that there can be no moral conflicts involving strong oughts, almost as if the idea that conflicts must be limited to weak oughts while strong oughts remain conflict free follows at once from a mere recognition of that distinction.

An example of a writer who seems simply to assume the idea is Searle, who suggests in [1980, pp. 248–249] that our moral reasoning can best be

represented through the postulation of a variety of deontic operators $\bigcirc_1, \bigcirc_2, \dots, \bigcirc_n$ to represent the weak obligations deriving from different sources, together with the special deontic operator \bigcirc^* representing the strong, all things considered ought. Having introduced this distinction, Searle then denies the validity of statements like $\bigcirc_1 A \supset \bigcirc_2 \neg A$ and $\bigcirc_1 A \supset \bigcirc^* \neg A$. That is, he denies, in the first case, that the fact that some action is obligatory according to one source of obligation entails that the opposite action cannot be obligatory according to another, thus allowing for conflicts among different weak obligations. And in the second case, he denies that the fact that some action is obligatory according to some particular source of obligation entails that the opposite cannot be what the agent ultimately ought to do, thus allowing for conflicts between weak obligations and the strong, all things considered moral ought. At this point, however, Searle simply notes—without comment or argument—that the statement $\bigcirc^* A \supset \neg \bigcirc^* \neg A$ should be classified as valid, ruling out the possibility of a conflict among strong oughts by asserting “that if one ought to do some particular thing, all things considered, then it is not the case that, all things considered, one ought not to do that very thing.” Searle’s paper, which is typical of the literature, thus contains a clear statement of the position that weak but not strong oughts should allow for conflicts, but no real argument for the position; it is simply taken for granted.

In fact, I can find only two writers, Thomson and Foot, who actually seem to argue for this position on the basis of conceptual considerations related to the nature of the distinction between weak and strong oughts. Thomson is, again, happy to admit that there might be conflicts among weak moral reasons for action, such as commitments or obligations, but rules out conflicts among strong oughts.

It should be stressed that what is an odd idea is that “I ought to give C a banana” and “I ought to give D a banana” are both true [even though I have only one banana to give, and cannot give it to both C and D]. There is no oddity in the idea that “I am committed to C to giving C a banana” and “I am committed to D to giving D a banana” are both true. Similarly for the ordinary English expressions “obligation” and duty”. [Thomson, 1990, p. 83]

To support the idea that there can be no conflicts among strong moral oughts, she sets out two arguments, the first of which appears in the passage immediately preceding that just cited.

Some philosophers have canvassed the idea in recent years that it can be the case that I ought to do alpha and ought to do beta despite the fact that I cannot do both alpha and beta. Should we agree with them? It is an odd idea. I will certainly feel you have been unhelpful if when I tell you about my predicament,

and ask what I ought to do, you tell me “Well, as a matter of fact, you ought to give C a banana and you ought to give D a banana.” I just told you I have only one banana. [Thomson, 1990, p. 83]

To this argument, I can offer two related replies. First, to say that your response to Thomson’s question is unhelpful—or odd, or in some other way inappropriate—is not necessarily to say that it is incorrect, as we know from the theory of conversational implicature. It could easily be that your response provides a correct statement of the moral facts, which, unfortunately, are not particularly helpful in this case. And second, it may be that the reason your response appears to be unhelpful, or odd, is that the response suggests a misunderstanding of the natural sense in which a question like Thomson’s would be asked. Here it is useful to recall Williams’s distinction between what might be called the *moral* and the *deliberative* senses of the word “ought”—between, that is, “the *ought* that occurs in statements of moral principle, and in the sorts of moral judgments about particular situations that we have been considering, [and] the *ought* that occurs in the deliberative question ‘what ought I to do?’ and in answers to this question, given by myself or another” [1965, p. 184].

As Williams points out, it is often natural to ask this kind of deliberative question even when all the moral facts are in place and agreed upon. Suppose, for example, that I face a very weak but nonetheless clear all things considered moral ought enjoining one action, but also have a compelling but largely self-interested reason for performing a conflicting action. In such a situation, it would be natural for me to ask a friend what I ought to do, where the point of this deliberative question would be to solicit help in balancing the weak moral demand against what is, in this case, a much stronger demand based on my own self-interest. It would be unhelpful, and odd, for my friend to respond to this question simply by reiterating the moral demand—the moral ought—which is presumably already known and taken as a premise of the question. The fact that such a response would be unhelpful, however, does not mean that it would be incorrect as a statement of the moral facts, only that I am not asking to be reminded of the moral facts. I am asking, instead, a deliberative question about what I ought to do when a number of considerations, including the moral facts, are taken into account; to suppose that the moral facts alone determine the answer to this deliberative question is to commit oneself to what Williams calls the “necessary supremacy of the moral.”

A similar point can be made about Thomson’s question regarding what she ought to do with the banana. Even though it might be unhelpful for you to respond to this question simply by reiterating the moral facts—that she ought to give the banana to C, but that she also ought to give the banana to D—this does not necessarily mean that your statement of the moral facts is incorrect. It may mean only that the question is naturally interpreted as

deliberative, taking the moral predicament as a premise and asking what she should do now that she has found herself in such a predicament. In this case, a more helpful reply might have the form: “Well, you’re stuck; what you ought to do now is flip a coin, give the banana to the winner, apologize to the loser, and offer him the very next banana you come across.” Of course, it is just possible to regard this last bit of advice as itself a dictate of morality, as if morality itself specifies the appropriate behavior in the face of the conflicting oughts, so that, in a sense, there is no real moral conflict. But it is also possible to imagine that morality generates but does not resolve the conflict, and that the advice is intended simply as a deliberative, practical suggestion as to the best thing to do once one has landed in such a moral predicament.

So much, then, for Thomson’s first argument in support of the view that there can be no conflicts among strong moral oughts. Her second argument occurs a few pages later, and rests on a claim about English usage.

I think myself that it was not merely odd but patently incorrect to think that “I ought to give C a banana” and “I ought to give D a banana” can both be true compatibly with my having only one banana; I think we simply do not use the English word “ought” in such a way that that is so. In any case, I will not. I will throughout so use “ought” that it cannot be the case that I ought to do alpha and ought to do beta where I cannot do both alpha and beta. [Thomson, 1990, p. 83]

This argument, however, can be disposed of quickly. It is, of course, always fair to stipulate that a term will be used in some particular way for one’s own local purposes, but as a claim about the English language, the idea that we do not use the word “ought” in a way that allows for conflicts is surely incorrect. There seems to be nothing especially odd to the ear about a statement like: “Oh dear, I promised Smith to have my paper finished by the end of the week, so I *should* stay home this evening and work, but Jones will be so disappointed if I don’t go to his party—I really *ought* to be there.” Even many of those authors who ultimately deny the possibility of moral conflict concede that English usage suggests otherwise, such as Foot, who writes on the basis of considerations about ordinary English, that “it may seem surprising that anyone should ever have denied that I can have an obligation to do *a* and an obligation not to do *a*, or that I ought to do *a* and ought not to do it” [1983, p. 254]. And as we have seen, other writers, such as Gowans, feel that the term “ought” itself is most naturally taken to refer to weak moral oughts, which allow conflicts, rather than our strong moral oughts, which are often taken to express what we “must” do.

Let us now turn to Foot’s own argument [1983, pp. 254–257] for the idea that weak, but not strong, oughts can allow for moral conflicts; the argument is based on the nature of the distinction she draws between weak and strong oughts, or, as she calls them, “type 1” and “type 2” oughts respectively. In Foot’s view, type 1 oughts are analogous to engagements—arrangements or

commitments that we might make to perform some action. Just as we can easily find ourselves with conflicting engagements, it is equally possible to face a conflict among type 1 oughts. Indeed, the possibility of conflict involving type 1 oughts, commitments or obligations, is so plain, Foot claims, that any resistance to the idea must be due largely to a confusion of type 1 oughts with type 2 oughts, which do not allow for conflicts, and which she introduces in the following passage:

What is a type 2 ought statement? What is it that makes 'ought a '(2) inconsistent with 'ought $\neg a$ '(2), although 'ought a '(1) is consistent with 'ought $\neg a$ '(1)?

The explanation is that type 2 ought statements tell us *the right* thing to do, and that this means the thing that is *best* morally speaking. . . . It is implied that for one for whom moral considerations are reasons to act, there are better moral reasons for doing this action than for any other. As this cannot be true both of a and of $\neg a$, 'ought a '(2) is inconsistent with 'ought $\neg a$ '(2). 'Ought a '(2) is not, however inconsistent with 'ought $\neg a$ '(1). I can have reason not to do something and yet have better reason to do it than I have to do anything else. [Foot, 1983, p. 256]

Given this way of defining type 2, or strong, oughts, Foot is correct in claiming that they cannot allow conflicts. Let a and b be conflicting actions, which cannot both be performed. Unless the preference ordering among reasons for acting is circular, there cannot be better reasons for performing a than for performing b , and also better reasons for performing b than for performing a . If 'ought x ' in the type 2 sense is taken to mean that there are better reasons for performing the action x than for performing any conflicting action, it follows that we cannot accept both 'ought a ' and 'ought b '. And what if the reasons for performing a and b are either incommensurable or identical in strength? In that case, since neither reason is actually better than the other, it follows that we can accept neither 'ought a ' nor 'ought b ', and conflict is avoided.

Foot's definition, then, does indeed entail that type 2 oughts do not allow for moral conflicts, but here we must object to the definition itself, on the grounds that it yields intuitively incorrect results. In fact, Foot's definition of type 2 oughts in terms of the best moral reasons—reasons that are better than those for performing any conflicting action—is essentially equivalent to Brink's suggestion, considered earlier, of defining all things considered oughts in terms of those *prima facie* oughts that are not only undefeated themselves, but also defeat all competitors—that is, the overriding oughts.

Since the two definitions are equivalent, they are subject to similar problems, as we can see by returning to the situation in which I have promised the two twins to have dinner with each of them tonight, but cannot have dinner with both, so that the two actions are inconsistent. In this situation, there is a reason, my promise, for performing each of these

two conflicting actions, but if we suppose that neither of these reasons is actually better than the other, then according to Foot's definition, neither action is something that I ought to do in the type 2 sense. The definition is thus successful in avoiding conflict, but just as we saw earlier, it also fails to provide us with any mechanism for drawing the intuitively desirable result that, given my promises, I ought at least to have dinner with one twin or the other; lacking further elaboration, there is nothing in Foot's account to tell us that this disjunctive action is supported by some best reason.

Of course, as we saw earlier, it is possible to avoid all things considered conflicts while at the same time generating the appropriate disjunctive oughts by adopting a different approach—basically, the disjunctive account as described here. Suppose, then, that we repair the technical error in Foot's definition by adopting a treatment more along the lines of our disjunctive account—stipulating, roughly, that an action is what one ought to do in the type 2 sense whenever it is supported, not by some reason that actually defeats all others, but instead, by each maximal consistent set of reasons that are not themselves defeated. In that case, the definition of type 2 oughts would yield intuitively acceptable results, and it would indeed follow from this definition that type 2 oughts cannot allow for conflicts. But what does this tell us, exactly? If we pattern the definition of type 2 oughts after something like the disjunctive account, they will be conflict free; but it is equally true that if we pattern the definition of type 2 oughts after something like the conflict account, they will allow for conflicts. Once we accept the technical emendation, Foot's paper can then be taken to show that there is a coherent way of defining a conflict-free notion of strong oughts, in addition to a notion that allows for conflicts; but the paper does not present an argument for adopting either of these two notions in preference to the other.

3.4 *Objections based on oughts as moral forces*

The final argument that I will consider here for rejecting the idea of conflicts among all things considered oughts is based on what is, frankly, a metaphor—the metaphor of *prima facie* oughts as moral forces, analogous in some ways to physical forces, and of all things considered oughts as the outcomes resulting from interactions among the various moral forces that are active in some particular situation. This metaphor can be found already in Ross's work:

Any act that we do contains various elements in virtue of which it falls under various categories. In virtue of being the breaking of a promise, for instance, it tends to be wrong; in virtue of being an instance of relieving distress it tends to be right. Tendency to be one's duty may be called a *parti-resultant* attribute, one which belongs to an act in virtue of some one component in its nature. *Being* one's duty is a *toti-resultant* attribute, one which belongs to an act in virtue of its whole nature and of nothing less than this.

Another instance of the same distinction may be found in the operation of natural laws. *Qua* subject to the force of gravitation towards some other body, each body tends to move in a particular direction with a particular velocity; but its actual movement depends on *all* the forces to which it is subject. [Ross, 1930, pp. 28–29]

As Ross himself suggests, the moral forces metaphor leads naturally to an interpretation of prima facie oughts as *ceteris paribus*—or “other things being equal”—moral statements, telling us what all things considered oughts would hold in the absence of any other moral forces; here, for example, is his commentary on the prima facie duty of returning good for good:

What I maintain is that an act in which good is returned for good is recognized as *specially* binding on us just because it is of that character, and that *ceteris paribus* any one would think it his duty to help his benefactors rather than his enemies, if he could not do both; just as it is generally recognized that *ceteris paribus* we should pay our debts rather than give our money in charity, when we cannot do both. [Ross, 1930, p. 30]

More exactly, the *ceteris paribus* interpretation of prima facie oughts holds that a statement of the form “Under the circumstances *A*, it ought prima facie to be that *B*” should be taken to mean that, whenever the circumstances *A* occur, then other things being equal—that is, in the absence of any additional moral forces—it ought all things considered to be that *B*.

Of course, it will rarely occur in any real situation that additional moral forces are entirely absent, that only a single moral consideration is relevant, but this is exactly what one would expect from the analogy between moral and physical forces. Even though some physical law might imply that a certain object would behave in a particular way in the absence of any additional physical forces, it is impossible to determine the exact behavior of that object in the real world until all the various forces actually influencing it are taken into account: the gravitational attraction between one object and another might imply a particular trajectory, for example, but the actual trajectory could then be influenced by friction, collisions, further gravitational attractions to other objects, and so on. In just the same way, Ross thought, one must take account of all the various moral forces to which an action is subject before arriving at an all things considered evaluation. The fact that some action—say, repaying a debt rather than donating to charity—is an instance of returning good for good tells us only that this action is what the agent ought to do in the absence of further moral forces. But in the usual case, there will be further moral forces, further prima facie oughts, to contend with: perhaps the agent has promised money to the charity, so that the prima facie ought concerning promising comes into play, or perhaps the charity is exceptionally benevo-

lent and efficient in its use of contributions, so that the prima facie ought of beneficence is relevant.

The metaphor of prima facie oughts as moral forces, with all things considered oughts as the resulting outcomes, and also the related interpretation of prima facie ought statements as ceteris paribus moral laws are both useful, I think. They are, in any case, pervasive, lying just below the surface in several discussions of moral reasoning, and have recently been defended explicitly and in detail both by Brink [1994, pp. 216–220] and by Paul Pietroski [1993]. What I do not see, however, is how—as both Brink and Pietroski suggest—accepting either the moral forces metaphor or the ceteris paribus interpretation of prima facie oughts provides a reason for rejecting the possibility of all things considered moral conflicts.¹⁵

To begin with, simply interpreting prima facie oughts as ceteris paribus moral laws does not tell us anything at all about the kind of conclusions we should expect to be derivable from a body of prima facie oughts as premises, and in particular, whether these conclusions can or cannot allow for conflicts; there is no generally accepted organon for reasoning with ceteris paribus laws, moral or otherwise.¹⁶ To say that it ought, ceteris paribus, to be the case that *B* is, again, to say only that it ought to be the case that *B* other things being equal, in the absence of any further moral forces. But since we so rarely encounter a situation in which only a single moral force comes into play, what we really need to know is how the various moral forces that are operative in a given situation interact with each other to determine the resulting outcomes—how the various prima facie oughts bearing on a situation interact to determine the resulting all things considered oughts.

The principles specifying the way in which input forces interact to determine the resulting outcomes in a force theory are sometimes known as “composition principles.” Of course, everyone’s favorite example of composition principles at work is found in classical physics, where both the input forces acting on an object and its resulting output behavior are represented as vectors, and the output is calculated from the inputs through simple vector arithmetic. The analogy with classical physics is followed rather closely by Brink, who speaks of the process through which prima facie moral forces determine all things considered oughts as “moral factor addition”:

It is not essential to the factor addition model that we always be able to assign precise numerical values to the various moral forces present in a situation. What is important is that the moral status of an act *sans phrase* results from adding the moral forces, positive and negative, contributed by the various morally relevant factors; the act with the highest moral total is all-things-considered obligatory. [Brink, 1994, p. 217]

But it is not necessary to follow classical physics so rigidly. Force composition principles from some of the special sciences—genetics, evolutionary

biology, economics, psychology—can differ substantially from those of classical physics. Often, in these special sciences, the composition principles involved are qualitative rather than quantitative, or statistical rather than deterministic; at times one force is allowed to override, or trump, another, rather than simply modifying its effect. And it is not unreasonable to expect that the principles governing the interactions among moral forces may be even more complicated.

What I would like to suggest is that both of the two accounts presented here—both the conflict account and the disjunctive account—can sensibly be regarded as providing exactly what is needed: an approximation, at least, of a set of appropriate composition principles for moral forces. Both accounts provide a precise method, subject to reasonable logical and intuitive constraints, for calculating all things considered oughts from a background context of *prima facie* oughts. If *prima facie* oughts are thought of as input moral forces and all things considered oughts as the resulting outputs, then both accounts can be seen as encoding principles through which outputs are determined by the input forces. And if this is so—if both the conflict and disjunctive accounts can sensibly be thought of as providing force composition principles—then it is hard to see how the moral force metaphor could be used as a basis for favoring either of these two accounts over the other.

It is, of course, possible to argue that, for some reason, it would be unnatural to interpret the conflict account as specifying composition principles for moral forces. Although one might happily allow conflicting moral forces as inputs, perhaps there is something odd about the idea of force composition principles that allow conflicts even among their outputs. Pietroski suggests as much in [1993, pp. 502–503]. He considers a situation in which some agent, Morty, is subject to conflicting *prima facie* oughts: first, to be at the train station on time to meet a friend, and second, to help a child he comes across on the way, which would then cause him to be late for the train. This situation is compared to a case in evolutionary biology in which some population is subject to conflicting *ceteris paribus* laws: one predicting drift in favor of some trait *T*, and another predicting selection in favor of a competitor trait *T'*. And here, Pietroski claims, the question of what Morty actually ought to do, in the all things considered sense, is analogous to the question concerning the direction in which the population will actually evolve.

Saying that Morty ought_{act} [that is, actually ought] to go to the station *and* that Morty ought_{act} to help the child makes no more sense than saying that the population will actually evolve in the direction of *T* *and* that it will actually evolve in the direction of *T'*. The population will evolve in the direction of *T*, or it will evolve in the direction of *T'*, or perhaps it will not evolve at all. . . . But the population cannot evolve in both directions. Nor can a ball simultaneously pushed north and pushed south move north and move south. Similarly, Morty

ought_{act} to be at the station on time, or (exclusively) he ought_{act} to help the child. But it is not that case that he ought to do both. [Pietroski, 1993, p. 503]

Is this a reasonable analogy? I agree, of course, that the ball will not simultaneously move north and south, and that the population will not evolve in both directions at once. These things simply cannot happen, and so any set of force composition principles predicting that they will happen would have to be rejected at once, on the grounds of predicting the impossible. But is it, in the same way, impossible to imagine that an agent might be subject to conflicting all things considered oughts, so that a theory predicting that he is would likewise have to be rejected? Again, this question—whether there can be conflicting all things considered oughts—is the very point at issue, and an answer cannot be taken for granted as a basis for argument.

Instead of physics or evolutionary biology, let us consider a different analogue for the behavior of moral forces. Suppose I have just taken a new job at, say, the Acme Corporation, to begin on Monday. Imagine that the corporate offices for Acme are located in a rural setting, so that the only practical way to getting to work is to drive there and park in the corporate parking garage. During my orientation session, the various rules governing Acme employees are explained to me. I learn that, in order to park in the corporate garage, it will be necessary to display an official Acme parking decal, which is to be mailed to my home address as part of my Employment Packet. If I park in the garage without displaying the decal, I will receive a ticket, resulting in a fine. I also learn that I will be docked pay for each day of work missed: in effect, another fine. In virtue of my employment at the Acme Corporation, I am therefore subject to certain forces that, if not moral, are at least normative, in the sense that I will receive a sanction—in this case, a fine—if I fail to act appropriately. One of these normative forces compels me to report to work each day; another compels me not to park in the corporate garage without displaying a parking decal.

Now suppose that Monday arrives, but that my Employment Packet, containing the parking decal, has not yet come in the mail. What should I do? If I fail to report to work, I will receive a fine; but going to work requires parking in the corporate garage, and I will also receive a fine if I park there without displaying a decal. At this point, it may seem reasonable for me to get in touch with Acme. So imagine that I call the Parking Coordinator to explain my situation. He understands the problem, and is sympathetic, but claims that there is little he can do: his job is simply to enforce the parking regulations, which require him to ticket any car parked in the garage without a decal. Imagine that I then call the only other relevant Acme official, the Personnel Director, who is likewise sympathetic but unable to help: his job is simply to enforce the personnel regulations, which require him to fine an employee for each missed day of work.

In this simple example, there is a clear and concrete criterion for determining whether I have violated a norm of the Acme Corporation: I receive the relevant fine. Apparently, I am now destined to receive some fine or another, either for parking without a decal or for missing a day of work. It is therefore natural to conclude that the normative forces to which I am subject in virtue of my employment at the Acme Corporation are organized in such a way that, under the circumstances in which I find myself, I simply cannot avoid violating some norm. And perhaps the composition principles governing moral forces work in exactly the same way, so that at times—even after balancing the relevant forces to the best of our ability, and filing the appropriate appeals—we may nevertheless be faced with a situation in which we cannot avoid violating some all things considered moral ought.

One might, of course, object to this analogy by arguing that there is something deeply flawed either with the employment regulations issued by the Acme Corporation or with their administration: surely there ought to be some individual outranking both the Parking and Personnel Directors, with the authority to adjudicate between the relevant rules in a case like mine, so that a normative conflict is avoided. Maybe so, but not all normative systems are organized and administered as we would hope. Perhaps Acme is just badly run. Or the case could be modified by supposing that the parking garage is operated by an entirely independent organization—Consolidated Parking, say. In this new situation, I would then be violating a Consolidated regulation if I go to work and an Acme regulation if I do not; but since these two organizations are independent, there would be no reason at all to expect any degree of coordination in the formulation or administration of their rules, or any higher authority to whom I could appeal a conflict.

Alternatively, one might respond to the analogy by arguing that the rules of morality must be better organized than the employment regulations issued by the Acme Corporation—and in particular that, although it may indeed be possible for employees to face certain normative conflicts involving the Acme regulations, there can be no conflicts regarding all things considered moral oughts. What a proponent of this view would need to establish is that, because of the special nature of morality among the variety of normative systems, moral rules—unlike employment regulations, for example—must either be so carefully qualified that the antecedents of any two rules with conflicting consequents can never both apply in the same situation, or else that the preference ordering on moral rules must be strongly connected, so that in case of conflict, one of the two rules will be given precedence. Again, this is an entirely coherent position, but it requires substantive moral argument for its justification. Until such an argument is provided, it is hard to see what reason there could be for ruling out the possibility of moral conflicts, or for preferring the disjunctive account to the conflict account.

A Appendix

The purpose of this appendix is to establish two results mentioned in Section 2.3 of the text: first, that the logic presented here as the conflict account is a conservative generalization of that defined by van Fraassen [1973], and second, that both the conflict and disjunctive accounts can be seen to agree with standard deontic logic, and therefore with each other, when applied to a consistent set of background oughts. We present these results in turn.

A.1 A comparison with van Fraassen's logic

In van Fraassen's account, ought statements are derived from a background set of imperatives, which can be identified with our *prima facie* oughts, but there are two ways in which his framework is more restrictive than that of the present paper: first, it deals only with unconditional ought statements and imperatives, and second, it does not allow a priority ordering on the background set of imperatives. We show here that when the present framework is subject to similar restrictions—that is, when we limit our consideration to unconditional statements of the form $!(B)$ and $\bigcirc(B)$, and when we assume that the priority ordering on imperatives is empty—then the results delivered by the conflict account coincide with those of van Fraassen. Just as a reminder, it is worth noting that, even though we focus here on unconditional imperatives of the form $!(B)$, we continue to abide by our present convention that these unconditional imperatives are officially defined as conditional imperatives that happen to be contingent on the trivial proposition \top . The unconditional imperative $!(B)$ is therefore identified with the imperative $!(B/\top)$, and so the consequent notation makes sense: if i is the imperative $!(B)$, then $\text{Consequent}[i]$ is the statement B .

As van Fraassen presents it, his account relies formally on a notion of *score*. Where v is an ordinary model of the underlying propositional language—that is, a simple valuation mapping sentence letters into truth values—the score of the valuation v , relative to a set \mathcal{I} of imperatives, is defined as the particular subset of imperatives from \mathcal{I} that are satisfied by v . As usual, we let \models represent the ordinary satisfaction relation between models and propositional formulas, so that $v \models A$ tells us that the formula A is satisfied by the model v . And as before, we suppose that an imperative i is satisfied by v whenever its consequent is itself satisfied—that is, whenever $v \models \text{Consequent}[i]$.

The score of a model v relative to a set \mathcal{I} of imperatives can therefore be defined as follows:

$$\text{score}_{\mathcal{I}}(v) = \{i \in \mathcal{I} : v \models \text{Consequent}[i]\}.$$

In this framework, we now let $|B|$ represent the ordinary model class of the formula B , the set of models that satisfies B , and where \mathcal{F} is a set of formulas, $|\mathcal{F}|$ is the intersection of the model classes of the formulas it contains: formally, $|B| = \{v : v \models B\}$ and $|\mathcal{F}| = \bigcap\{|B| : B \in \mathcal{F}\}$. Van Fraassen's notion of deontic consequence, which we represent here as the relation \sim_F , is then defined as follows.

- Let \mathcal{I} be some set of unconditional imperatives. Then $\mathcal{I} \sim_F \bigcirc(B)$ if and only if there is a model $v_1 \in |B|$ for which there is no model $v_2 \in |\neg B|$ such that $\text{score}_{\mathcal{I}}(v_1) \subseteq \text{score}_{\mathcal{I}}(v_2)$.

The idea is that B ought to be the case, given the background set \mathcal{I} of imperatives, just in case the truth of B is a necessary condition for achieving some maximal score.

Let us now consider the shape of the conflict account under the current restrictions. It is easy enough to see that, when we restrict our consideration to unconditional oughts, with no priority ordering, then every prima facie ought from the background context will be binding under any circumstances whatsoever. More exactly, given a background context $\langle \mathcal{I}, \emptyset \rangle$, in which \mathcal{I} contains only unconditional prima facie oughts and \emptyset represents the empty ordering, we will have $\text{Binding}_{\langle \mathcal{I}, \emptyset \rangle}(A) = \mathcal{I}$, so that $\text{Consequent}[\text{Binding}_{\langle \mathcal{I}, \emptyset \rangle}(A)] = \text{Consequent}[\mathcal{I}]$. This observation leads at once to the following fact, which records a special case of the general evaluation rule for the conflict account that is applicable under the current restrictions.

Fact 1 Let $\langle \mathcal{I}, \emptyset \rangle$ be a background context in which \mathcal{I} contains only unconditional prima facie oughts. Then $\langle \mathcal{I}, \emptyset \rangle \sim_C \circ (B)$ if and only if $\mathcal{F} \vdash B$ for some maximal consistent subset \mathcal{F} of the set $\text{Consequent}[\mathcal{I}]$.

The coincidence between van Fraassen's logic and this restricted case of the conflict account can then be stated as follows.

Fact 2 Let $\langle \mathcal{I}, \emptyset \rangle$ be a background context in which \mathcal{I} contains only unconditional prima facie oughts. Then $\langle \mathcal{I}, \emptyset \rangle \sim_C \circ (B)$ if and only if $\mathcal{I} \sim_F \circ (B)$.

Proof First, suppose that $\mathcal{I} \sim_F \circ (B)$. Then there is a model $v_1 \in |A|$ for which there is no model $v_2 \in |\neg A|$ such that $\text{score}_{\mathcal{I}}(v_1) \subseteq \text{score}_{\mathcal{I}}(v_2)$. Taking $\text{Th}(v)$ to stand as usual for the set of formulas true in a model—that is, $\text{Th}(v) = \{A : v \models A\}$ —we let

$$\mathcal{G} = \text{Th}(v_1) \cap \text{Consequent}[\mathcal{I}].$$

Clearly, \mathcal{G} is consistent and a subset of $\text{Consequent}[\mathcal{I}]$; and it is clear also that $\text{score}_{\mathcal{I}}(v) = \text{score}_{\mathcal{I}}(v')$ for any two models $v, v' \in |\mathcal{G}|$. To see that $\mathcal{G} \vdash B$, suppose otherwise: then there exists a model $v_2 \in |\mathcal{G}| \cap |\neg B|$; but in that case we have $\text{score}_{\mathcal{I}}(v_2) = \text{score}_{\mathcal{I}}(v_1)$, contrary to the definition of \sim_F . The set \mathcal{G} is therefore a consistent subset of $\text{Consequent}[\mathcal{I}]$ such that $\mathcal{G} \vdash B$. Standard techniques then allow \mathcal{G} to be extended to a maximal consistent subset \mathcal{F} of $\text{Consequent}[\mathcal{I}]$ such that $\mathcal{F} \vdash B$, and so we have $\langle \mathcal{I}, \emptyset \rangle \sim_C \circ (B)$ by Fact 1.

Next, suppose that $\langle \mathcal{I}, \emptyset \rangle \sim_C \circ (B)$, so that there is some maximal consistent subset \mathcal{F} of $\text{Consequent}[\mathcal{I}]$ such that $\mathcal{F} \vdash B$. Since \mathcal{F} is consistent, and since $\mathcal{F} \vdash B$, we have some model $v_1 \in |\mathcal{F}| \subseteq |B|$; and then since \mathcal{F} is maximal, it is easy to see that there can be no $v_2 \in |\neg B|$ such that $\text{score}_{\mathcal{I}}(v_1) \subseteq \text{score}_{\mathcal{I}}(v_2)$. So $\mathcal{I} \sim_F \circ (B)$.

A.2 A comparison with standard deontic logic

Unlike van Fraassen's proposal, standard deontic logic is a species of modal logic, developed using the usual possible worlds techniques; accessible treatments can be found in most texts on modal logic, such as Chellas [1980].

Very briefly, standard deontic logic is the modal logic based on *standard deontic models*—structures of the form $\mathcal{M} = \langle W, f, v \rangle$, with W a set of possible worlds, v a modal valuation mapping sentence letters into sets of worlds at which they are thought of as true, and f a function mapping each world α into a set of worlds $f(\alpha)$, subject only to the constraint that this set of worlds should be nonempty: $f(\alpha) \neq \emptyset$. Where α is an individual world, $f(\alpha)$ can be thought of as the set of worlds that are ideal from the standpoint of α , those in which all the oughts in force at α are satisfied.

Following the usual pattern in modal logics, formulas are now assigned truth values relative to a pair consisting of a model \mathcal{M} and a world α from that model. For a sentence letter p , we have the evaluation rule

$$\mathcal{M}, \alpha \models p \text{ if and only if } \alpha \in v(p),$$

telling us simply that p is true at the world α if α is among the worlds assigned by v to p . The rules for the truth functional connectives mirror those of ordinary logic:

$$\begin{aligned} \mathcal{M}, \alpha \models A \wedge B & \text{ if and only if } \mathcal{M}, \alpha \models A \text{ and } \mathcal{M}, \alpha \models B, \\ \mathcal{M}, \alpha \models \neg A & \text{ if and only if it is not the case that } \mathcal{M}, \alpha \models A \end{aligned}$$

And the rule for the deontic operator \bigcirc follows the standard recipe

$$\mathcal{M}, \alpha \models \bigcirc(B) \text{ if and only if } \mathcal{M}, \beta \models B \text{ for each } \beta \in f(\alpha).$$

The idea is that $\bigcirc(B)$ holds at a world α in the model \mathcal{M} just in case B holds in all the worlds that are ideal from the standpoint of α .

As usual, we will say that $\mathcal{M} \models A$ just in case A holds at each world in \mathcal{M} —just in case, that is, $\mathcal{M}, \alpha \models A$ for each world α from the model. Where Γ is a set of formulas, we will say that $\mathcal{M} \models \Gamma$ whenever $\mathcal{M} \models A$ for each formula A from Γ . And we will say that the set of formulas Γ semantically entails the formula A according to standard deontic logic—written $\Gamma \models_{SDLA} A$ —just in case $\mathcal{M} \models \Gamma$ implies $\mathcal{M} \models A$ for each standard deontic model \mathcal{M} .

In order to establish the appropriate connection between standard deontic logic and the accounts presented here, we must first impose a number of restrictions on both formalisms, beginning with those set out earlier: we consider only those contexts of the form $\langle \mathcal{I}, \emptyset \rangle$ in which \mathcal{I} contains only unconditional prima facie oughts, and in which the priority ordering among prima facie oughts is empty. Since standard deontic logic allows nested deontic operators, while the systems presented here do not, we must, next, explicitly restrict ourselves only to the non-nested fragment of standard deontic logic. And finally, since standard deontic logic contains only a single deontic operator while the systems presented here contain two, distinguishing prima facie from all things considered oughts, our comparison will depend on a mapping that collapses these two deontic operators into one. Formally, where \mathcal{I} is a set of unconditional prima facie oughts, let us define the set $\mathcal{I}^* = \{\bigcirc(B) : !(B) \in \mathcal{I}\}$, so that \mathcal{I}^* represents, in a sense, the interpretation of \mathcal{I} in standard deontic logic.

With these restrictions in place, we can now turn to our primary observation—that both the conflict and disjunctive accounts coincide with standard deontic logic when applied to a consistent set of background oughts. A set \mathcal{I} of unconditional prima facie oughts will be defined as *consistent* whenever the set $Consequent[\mathcal{I}]$ is itself consistent—whenever, that is, the set of oughts \mathcal{I} is jointly satisfiable. As a preliminary observation, we note that, when \mathcal{I} is consistent in this sense, the set $Consequent[\mathcal{I}]$ will be its own unique maximal consistent subset. We are thus led to the following special case of the general evaluation rules for both conflict and

disjunctive oughts, which shows their coincidence when applied to consistent background contexts.

Fact 3 Let $\langle \mathcal{I}, \emptyset \rangle$ be a background context in which \mathcal{I} is a consistent set of unconditional prima facie oughts. Then we have both $\langle \mathcal{I}, \emptyset \rangle \sim_C \bigcirc(B)$ and $\langle \mathcal{I}, \emptyset \rangle \sim_D \bigcirc(B)$ if and only if $\text{Consequent}[\mathcal{I}] \vdash B$.

Because the conflict and disjunctive accounts coincide in this special case, where \mathcal{I} is consistent, we can in this case take $\langle \mathcal{I}, \emptyset \rangle \sim \bigcirc(B)$ to mean that $\bigcirc(B)$ follows from the context $\langle \mathcal{I}, \emptyset \rangle$ according to both the conflict or disjunctive accounts, indiscriminately. Our primary observation, establishing the agreement in case of consistency between these accounts and standard deontic logic, can therefore be stated as follows.

Fact 4 Let $\langle \mathcal{I}, \emptyset \rangle$ be a background context in which \mathcal{I} is a consistent set of unconditional prima facie oughts, and let B be an ordinary propositional formula. Then $\langle \mathcal{I}, \emptyset \rangle \sim \bigcirc(B)$ if and only if $\mathcal{I}^* \Vdash_{SDL} \bigcirc(B)$.

Proof We begin by supposing that $\mathcal{I}^* \Vdash_{SDL} \bigcirc(B)$ —that is, that $\mathcal{M} \models \mathcal{I}^*$ implies $\mathcal{M} \models \bigcirc(B)$ for each standard deontic model \mathcal{M} .

In order to show that $\langle \mathcal{I}, \emptyset \rangle \sim \bigcirc(B)$, we begin by constructing a particular standard deontic model $\mathcal{M}_{\mathcal{I}} = \langle W, f, v \rangle$, whose components are defined as follows. (1) W contains the set of models, ordinary valuations, for the underlying propositional language. This stipulation is subtle and potentially confusing, since the same objects— α, β, γ , and so on—now play two roles: they are both models or valuations for the underlying propositional language and also possible worlds in a particular model for the deontic language. (2) f is the constant function mapping each possible world α into the set $|\text{Consequent}[\mathcal{I}]|$. What this means is that f associates each world, each ordinary valuation, with the set of ordinary valuations that satisfy all the prima facie oughts from \mathcal{I} . (3) v is the modal valuation defined by taking $v(p) = \{\alpha : \alpha \models p\}$. What this means is that v maps each sentence letter p into the set of possible worlds that, considered now as ordinary propositional models, assign the value of truth to p .

At this point, we need to establish two preliminary facts.

First, to show that $\mathcal{M}_{\mathcal{I}}$ is, in fact, a standard deontic model, we need to guarantee that $f(\alpha) \neq \emptyset$, but this is trivial: $f(\alpha) = |\text{Consequent}[\mathcal{I}]|$ is simply the set of ordinary propositional models satisfying the set of sentences $\text{Consequent}[\mathcal{I}]$, but by hypothesis, $\text{Consequent}[\mathcal{I}]$ is a consistent set, and so we know that it must have at least one model.

Second, we need to show that, for each ordinary propositional formula B , we have

$$(*) \quad \mathcal{M}_{\mathcal{I}}, \beta \models B \text{ if and only if } \beta \models B,$$

telling us that B is satisfied by the world β in the model $\mathcal{M}_{\mathcal{I}}$ just in case β , now considered as an ordinary propositional valuation, assigns the value truth to B —or, put another way, that $\mathcal{M}_{\mathcal{I}}, \beta \models B$ if and only if $\beta \in |B|$, where again, $|B| = \{\beta : \beta \models B\}$ is the set of ordinary propositional valuations satisfying B . This fact can be established by induction on the complexity of the formula B . The base case is guaranteed by the definition of v in (3) above; the inductive step is straightforward.

We can now proceed with our main argument. Since $\mathcal{I}^* \Vdash_{SDL} \bigcirc(B)$, and since $\mathcal{M}_{\mathcal{I}}$ is a standard deontic model, we know that $\mathcal{M}_{\mathcal{I}} \models \bigcirc(B)$ whenever $\mathcal{M}_{\mathcal{I}} \models \mathcal{I}^*$. It is easy to verify that $\mathcal{M}_{\mathcal{I}} \models \mathcal{I}^*$, and so we can conclude that $\mathcal{M}_{\mathcal{I}} \models \bigcirc(B)$. According to the deontic evaluation rule, however, we have $\mathcal{M}_{\mathcal{I}} \models \bigcirc(B)$ only if $\mathcal{M}, \beta \models B$ for each $\beta \in f(\alpha)$. From (*) above, we can now conclude that $f(\alpha) \subseteq |B|$: each model belonging to $f(\alpha)$ —that is, each model satisfying *Consequent* $[\mathcal{I}]$ —is also a model of B . From this, it follows by the completeness theorem for ordinary propositional logic that *Consequent* $[\mathcal{I}] \vdash B$, from which we can conclude that $\langle \mathcal{I}, \emptyset \rangle \sim \bigcirc(B)$ by Fact 3.

The argument in the other direction is easier. If we suppose that $\langle \mathcal{I}, \emptyset \rangle \sim \bigcirc(B)$, we know by Fact 3 that *Consequent* $[\mathcal{I}] \vdash B$, so that $\{A_1, \dots, A_n\} \vdash B$ for some finite subset $\{A_1, \dots, A_n\}$ of *Consequent* $[\mathcal{I}]$. From this we can conclude that $\vdash (A_1 \wedge \dots \wedge A_n) \supset B$, by the deduction theorem for propositional logic. Since standard deontic logic is a normal modal logic, we therefore have $\Vdash_{SDL} (\bigcirc(A_1) \wedge \dots \wedge \bigcirc(A_n)) \supset \bigcirc(B)$, from which it follows that $\mathcal{I}^* \Vdash_{SDL} \bigcirc(B)$, since $\bigcirc(A_1), \dots, \bigcirc(A_n) \in \mathcal{I}^*$.

Acknowledgments

Michael Slote provided advice, encouragement, and a sense of perspective throughout the entire process of composition; Paul Pietroski offered further advice, further encouragement, and performed a sanity check on the final version of the paper. I am very grateful to them both.

Notes

¹ Although some writers, such as Michael Slote [1985], suggest that utilitarianism can allow for the possibility of moral conflicts, I do not know of any successful argument along these lines; I show in my [2001] that situations of the sort described by Slote, while anomalous, do not lead to moral conflicts in the sense defined here.

² A recent, sustained treatment of the logic of personal agency is presented by Nuel Belnap, Michael Perloff, and Ming Xu in their [2001]; the framework developed there is applied to the analysis of normative statements in my [2001], which includes an extensive discussion of certain differences between personal and impersonal ought statements.

³ The importance of symmetrical cases like this in reasoning about moral conflicts was first emphasized by Ruth Barcan Marcus [1980].

⁴ The point can be illustrated with a few examples. In his presidential address to the Aristotelian Society, D. D. Raphael [1975, pp. 11–12] argues that a choice between conflicting oughts must be regarded as resulting from a rational process, since it is possible to look back on the choice and view it as a mistake, an error; but he admits that he cannot actually find any rational criterion for evaluating such a choice. Later, in an appendix to this paper written after the address was delivered, Raphael explores the idea that the standards involved in justifying a choice among conflicting *prima facie* oughts might be those of

‘rhetorical’ reasoning, such as goes on in debate, and notably in judicial debate. . . . It is notorious that one cannot give a precise form to the alleged logic of rhetorical reasoning.

Nevertheless such debate and deliberation are not altogether irrational or non-rational.
[Raphael, 1975, p. 12d]

However, while it is certainly true that debate and judicial reasoning can be viewed as rational, it is equally true that, even after extensive debate, the opposing parties often fail to settle on a common conclusion; and of course, one influential view of judicial reasoning is that the law, at times, simply does not determine the result in a particular case: the decision is left to a judge, who must then actually extend the law to cover the new case. Thomas Nagel [1979, pp. 139–141] is likewise concerned with the problem of rational deliberation in situations in which different and conflicting values are brought to bear on a particular decision, but the method he sketches is designed to assure only that all the relevant considerations are attended to, not with the problem of actually reaching a decision based on these considerations—and of course, even assuring that the relevant considerations are heard provides no guarantee that a consensus will be reached. More recently, Henry Richardson [1990, p. 302] has argued that a conflict between norms should be resolved through a process of specifying the norms in such a way that the conflict is avoided; but then we face the problem of justifying one way of specifying the norms over another, to which Richardson responds only that a particular “specification is rationally defensible. . . so long as it enhances the mutual support among the set of norms found acceptable on reflection.” This may very well be correct, but defining a rationally defensible refinement of a set of norms in terms of what is found to be “acceptable on reflection” does not take us very far beyond the standpoint of Ross or Bradley, who define what is right in a particular situation as what would be perceived to be right by the practically wise person.

⁵ To say that a *prima facie* ought is defeated is not to say that it has no further moral impact at all—that it is, to echo Bernard Williams [1965], eliminated from the scene. A defeated *prima facie* ought may nevertheless justify feelings of compunction or regret, or even generate certain reparational oughts, such as the need to explain or apologize: in the present example, after missing a lunch date with my friend in order to rescue the child, I would certainly feel the need at least to explain the matter to my friend, and perhaps to apologize as well, in a way that I would not if we had never even had a lunch date. The problem of characterizing the moral force carried by defeated *prima facie* oughts is, however, a complicated matter that I will not try to address here.

⁶ A view similar to Brink’s was suggested earlier by Gilbert Harman, who presents his “good-reasons” analysis of ought statements as follows:

... to say that P ought to do D is to say that P has sufficient reasons to do D that are stronger than the reasons he has to do something else. If what you mean is that P morally ought to do D, you mean that P has sufficient *moral* reasons to do D that are stronger than the reasons he has to do anything else. [Harman, 1978, p. 118]

What Harman is telling us here is that oughts are to be defined in terms of reasons that are actually stronger than any conflicting reasons—that is, in term of overriding reasons—rather than simply in terms of reasons for which there are no stronger conflicting reasons, the undefeated reasons.

⁷ Conditional deontic logics of this variety include the theories presented by Bengt Hansson [1971], Dagfinn Føllesdal and Risto Hilpinen [1971], van Fraassen [1972], and David Lewis [1973]; a useful comparison of these various logics can be found in Lewis [1974].

⁸ Some preliminary work along these lines, relating nonmonotonic and deontic logics, can be found in my [1994] and [1997].

⁹ I am neglecting a number of familiar objections to the idea that the set of oughts should be closed under consequence, since these issues have no bearing on the present paper; some recent defenses of closure under consequence are set out by Brink [1994, pp. 234–235] and Thomson [1990, pp. 156–157].

¹⁰ Marcus [1980] refers to this principle as the rule of *factoring*.

¹¹ It was the desire to allow for some measure of agglomeration that led van Fraassen to move from the simple deontic logic presented in Section 6 of his [1973] to the more complicated system presented in Section 7. It was this desire also that lies behind my criticism in [1997] of proposals, such as that set out in Chapter 6 of Brian Chellas's [1980], to formalize deontic reasoning with conflicting oughts within the framework of weak, nonnormal modal logics.

¹² The reader will note that the rule of consistent agglomeration, like agglomeration itself, is formulated as a two-premise rule, while the rule of consistent consequent agglomeration requires an indefinite number of premises. The reason for this annoying difference in formulation is simply that consistent agglomeration can achieve the effect of a multiple-premise rule through iteration, while the multiple premises in consistent consequent agglomeration must be handled all at once. Imagine, for example, that we have derived the three all things considered oughts $\bigcirc(B_1)$, $\bigcirc(B_2)$, and $\bigcirc(B_3)$, from the three prima facie oughts $!(B_1)$, $!(B_2)$, and $!(B_3)$, and that the statement set $\{B_1, B_2, B_3\}$ is consistent. Then we can reach the agglomerated result $\bigcirc((B_1 \wedge B_2) \wedge B_3)$, through two application of consistent agglomeration, first combining $\bigcirc(B_1)$ and $\bigcirc(B_2)$ to yield $\bigcirc(B_1 \wedge B_2)$, and then combining $\bigcirc(B_1 \wedge B_2)$ with $\bigcirc(B_3)$ to reach $\bigcirc((B_1 \wedge B_2) \wedge B_3)$. In the case of consistent consequent agglomeration, however, this second step is blocked, since the formula $B_1 \wedge B_2$ is not the consequent of any prima facie oughts; given this restriction, the desired result $\bigcirc((B_1 \wedge B_2) \wedge B_3)$ can be reached only by agglomerating all three of our original ought formulas at once.

¹³ For the sake of tidiness, we note that this step in the argument requires that the relation \vdash_C should be closed under ordinary logical consequence, at least in the sense that, when we have both $\langle \mathcal{I}, \leq \rangle \vdash_C \phi$ and $\langle \mathcal{I}, \leq \rangle \vdash_C \phi \supset \psi$, we can then conclude that $\langle \mathcal{I}, \leq \rangle \vdash_C \psi$.

¹⁴ Brink refers to this principle as weak impermissibility because he formulates it as the principle that $\neg P(B)$ implies $\neg \bigcirc(B)$; the two formulations are equivalent, of course.

¹⁵ The idea that adopting the moral forces metaphor entails rejecting all things considered moral conflicts—that is, the implication between these two positions—is accepted, not only by those like Brink and Pietroski, who do adopt the metaphor and therefore reject moral conflicts, but also by some who wish to argue in the other direction. As an example, Jonathan Dancy [1993, pp. 102–103], who accepts the possibility of moral conflicts, argues against the moral forces metaphor on the grounds that it does not allow for this possibility.

¹⁶ Perhaps the most promising set of techniques for reasoning with *ceteris paribus* laws are those recently developed in the field of nonmonotonic logic. But this field contains a number of distinct approaches and theories, and as I have shown in my [1994], when prima facie oughts are formulated as *ceteris paribus* moral laws in some of the most prominent of these theories, it is then very natural to derive conflicting all things considered oughts as conclusions.

References

- Nuel Belnap, Michael Perloff, and Ming Xu. (2001) *Facing the Future: Agents and Choices in Our Indeterministic World*. Oxford University Press, 2001.
- David Brink. (1994) Moral conflict and its structure. *The Philosophical Review*, 103:215–247, 1994.
- Brian Chellas. (1980) *Modal Logic: An Introduction*. Cambridge University Press, 1980.
- Roderick Chisholm. (1964) The ethics of requirement. *American Philosophical Quarterly*, 1: 147–153, 1964.
- Earl Conee. (1982) Against moral dilemmas. *Philosophical Review*, 91:87–97, 1982. Reprinted in [Gowans, 1987b], pages 239–249; pagination refers to this version.
- Jonathan Dancy. (1993) *Moral Reasons*. Blackwell Publishers, 1993.

- Donald Davidson. (1970) How is weakness of the will possible? In Joel Feinberg, editor, *Moral Concepts*. Oxford University Press, 1970.
- Alan Donagan. (1984) Consistency in rationalist moral systems. *The Journal of Philosophy*, 81:291–309, 1984. Reprinted in [Gowans, 1987b], pages 271–290; pagination refers to this version.
- Alan Donagan. (1993) Moral dilemmas, genuine and spurious: a comparative anatomy. *Ethics*, 104:7–21, 1993.
- Dagfinn Føllesdal and Risto Hilpinen. (1971) Deontic logic: an introduction. In Risto Hilpinen, editor, *Deontic Logic: Introductory and Systematic Readings*, pages 1–35. D. Reidel Publishing Company, 1971.
- Philippa Foot. (1983) Moral realism and moral dilemma. *Journal of Philosophy*, 80:379–398, 1983. Reprinted in [Gowans, 1987b], pages 250–270; pagination refers to this version.
- Christopher Gowans. (1987a) Introduction: the debate on moral dilemmas. In Christopher Gowans, editor, *Moral Dilemmas*, pages 3–33. Oxford University Press, 1987.
- Christopher Gowans, editor. (1987b) *Moral Dilemmas*. Oxford University Press, 1987.
- Bengt Hansson. (1971) An analysis of some deontic logics. In Risto Hilpinen, editor, *Deontic Logic: Introductory and Systematic Readings*, pages 121–147. D. Reidel Publishing Company, 1971.
- Gilbert Harman. (1978) *The Nature of Morality: An Introduction to Ethics*. Oxford University Press, 1978.
- John Horty. (1994) Moral dilemmas and nonmonotonic logic. *Journal of Philosophical Logic*, 23:35–65, 1994.
- John Horty. (1997) Nonmonotonic foundations for deontic logic. In Donald Nute, editor, *Defeasible Deontic Logic*, pages 17–44. Kluwer Academic Publishers, 1997.
- John Horty. (2001) *Agency and Deontic Logic*. Oxford University Press, 2001.
- David Lewis. (1973) *Counterfactuals*. Oxford University Press, 1973.
- David Lewis. (1974) Semantic analyses for dyadic deontic logic. In Søren Stenlund, editor, *Logical Theory and Semantic Analysis*, pages 1–14. D. Reidel Publishing Company, 1974.
- Ruth Barcan Marcus. (1980) Moral dilemmas and consistency. *Journal of Philosophy*, 77: 121–136, 1980. Reprinted in [Gowans, 1987b], pages 188–204; pagination refers to this version.
- Thomas Nagel. (1979) The fragmentation of value. In Thomas Nagel, editor, *Mortal Questions*, pages 128–141. Cambridge University Press, 1979.
- Paul Pietroski. (1993) Prima facie obligations, ceteris paribus laws in moral theory. *Ethics*, 103: 489–515, 1993.
- D. D. Raphael. (1975) The standard of morals. *Proceedings of the Aristotelian Society*, 75: 1–12e, 1975.
- Henry Richardson. (1990) Specifying norms as a way to resolve concrete ethical problems. *Philosophy and Public Affairs*, 1990.
- W. D. Ross. (1930) *The Right and the Good*. Oxford University Press, 1930.
- Jean Paul Sartre. (1946) *L'Existentialisme est un Humanisme*. Nagel, 1946. Translated as “Existentialism is a Humanism” in *Existentialism from Dostoevsky to Sartre*, W. Kaufmann (ed.), Meridian Press, 1975.
- John Searle. (1980) Prima-facie obligations. In Zak van Straaten, editor, *Philosophical Subjects: Essays Presented to P. F. Strawson*, pages 238–259. Oxford University Press, 1980.
- Michael Slote. (1985) Utilitarianism, moral dilemmas, and moral cost. *American Philosophical Quarterly*, 22:161–168, 1985.
- Judith Jarvis Thomson. (1990) *The Realm of Rights*. Harvard University Press, 1990.
- Bas van Fraassen. (1972) The logic of conditional obligation. *The Journal of Philosophical Logic*, 72:417–438, 1972.

- Bas van Fraassen. (1973) Values and the heart's command. *The Journal of Philosophy*, 70:5–19, 1973.
- Bernard Williams. (1965) Ethical consistency. *Proceedings of the Aristotelian Society*, 39 (supplemental):103–124, 1965. A revised version appears in *Problems of the Self: Philosophical Papers 1956–1972*, Cambridge University Press, 1973, pages 166–186.