

Campbell Brown

Maximalism and the structure of acts

**Article (Accepted version)
(Refereed)**

Original citation:

Brown, Campbell (2017) *Maximalism and the structure of acts*. [Noûs](#). ISSN 0029-4624 (In Press)
DOI: [10.1111/nous.12181](https://doi.org/10.1111/nous.12181)

© 2017 Wiley Periodicals, Inc.

This version available at: <http://eprints.lse.ac.uk/88128/>
Available in LSE Research Online: June 2018

LSE has developed LSE Research Online so that users may access research output of the School. Copyright © and Moral Rights for the papers on this site are retained by the individual authors and/or other copyright owners. Users may download and/or print one copy of any article(s) in LSE Research Online to facilitate their private study or for non-commercial research. You may not engage in further distribution of the material or use it for any profit-making activities or any commercial gain. You may freely distribute the URL (<http://eprints.lse.ac.uk>) of the LSE Research Online website.

This document is the author's final accepted version of the journal article. There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

Maximalism and the Structure of Acts

Campbell Brown
London School of Economics
Final draft
Published in *Noûs*

DOI: 10.1111/nous.12181

Abstract. Suppose we believe that a property F is coextensive with moral permissibility. F may be, for example, the property of having the best consequences, if we are Consequentialists, or that of conforming to a universalisable maxim, if we are Kantians, and so on. A way to test this belief is to consider whether F behaves in the way that we expect permissibility to behave. An expectation we may have is that permissibility is “closed under implication”: any act that is implied by a permissible act must itself be permissible. Yet, in some cases, F might not be closed under implication. Maximalism has been proposed as a solution to this problem. A “maximal” act is one not implied by any other act. Maximalism restricts the claim that F is coextensive with permissibility to maximal acts only. A non-maximal act may be permissible without being F if it is implied by a maximal act that is F . The general aim of this paper is to investigate these issues by considering the formal structure of acts, or the “act-implication” relation. Discussions of Maximalism have tended to assume implicitly that acts have structure of some sort, but there has been little careful attention given to this structure. I aim to show that, by thinking about structure, we can provide a stronger defence of Maximalism.

1 Introduction

Some acts can be done by doing others.¹ You can drink whisky by drinking Lagavulin, visit Germany by visiting Munich, raise one of your arms by raising your left arm, and so on. When acts are related in this way, I shall say that the more specific act *implies* the less. Drinking Lagavulin implies drinking whisky, for example. As one sentence implies another when it is impossible for the first to be true without the second also being true, so one act implies another when it is impossible for the first to be done without the second also being done.²

This relation of implication between acts raises interesting moral questions. How might the moral properties of an act be related to those of an act it implies? One plausible principle seems to be that moral permissibility is “closed under implication”: any act implied by a permissible act must itself be permissible. Imagine an immigration official telling you that you’re permitted to enter Munich but not to enter Germany. That would be hard to comprehend. The official must have misspoken, you would think,

¹I say “act” rather than “action”, throughout this paper, solely because this seems the more common terminology in the relevant literature.

²In the case of acts, the intended grade of possibility must be stronger than logical possibility, since it is *logically* possible to visit Munich without visiting Germany, geographical truths not being logical truths. We might call it *practical* possibility. An act α is practically possible, for a given agent at a given time, if it is within the agent’s power at this time to do α . It is logically but not practically possible for me now to fly to the moon, for example.

or have a surprisingly poor grasp of geography. Or imagine a supermarket shopper who stands in the lane marked “Twelve Items or Less” while holding thirteen items. When informed that she is not allowed thirteen items in this lane, she responds: “But that is not what the sign says; all it says is that I am not allowed more than twelve.” Surely that is absurd. Having thirteen items *implies* having more than twelve; so you cannot be permitted to do the former without also being permitted to do the latter.

However, this principle might be problematic for some moral theories. Other properties which, according to these theories, are coextensive with moral permissibility might not be closed under implication. Consider, for example, Consequentialism. This theory holds that moral permissibility is coextensive with what we may call “optimality”, that is, the property of having consequences no worse than those of any alternative act.³ Some widely discussed examples suggest that optimality is not closed under implication.

These examples rest on an assumption about the consequences of an act. If an act can be done in various ways, then the consequences of doing it in one way may be better or worse than those of doing it in another way. Consider, for example, the act of driving to the shopping mall. This can be done safely or recklessly; sober or under the influence of alcohol; and so on. What, then, should we say are *the* consequences of such an act? A common assumption is that, among the ways of doing an act, one of these is special: it is the way that the act would in fact be done, if it were done at all. This is not to say, it should be stressed, that it is impossible to do the act in any other way. That something will in fact happen is compatible with the possibility of its not happening. I know that I will have breakfast tomorrow by eating a bagel, this being what I always do, but it is still possible for me to have something else for breakfast. There’s porridge in the cupboard that I *could* have instead, though I know that I won’t. The assumption, then, is that the relevant consequences of an act are those of doing the act in the way that it would be done.⁴ If I would drive to the shopping mall safely and soberly, then the consequences of my driving to the shopping mall are those of my driving safely and soberly. (This assumption is not beyond doubt. But I will not question it here. My interest is rather in the problem that emerges if we do grant the assumption.)

Given this assumption, examples in which optimality is not closed under implication are easy to construct. Consider two acts, α and β , each of which can be done in only one way. The “disjunction” of these acts, $\alpha \vee \beta$, can be done either by doing α or by doing β . (I’ll say more about act disjunctions below. I hope we can, for now, take them for granted.) Now suppose the following are true: α is optimal; β is not optimal; and $\alpha \vee \beta$ would be done by doing β . Then, since the relevant consequences of $\alpha \vee \beta$ are those of β , $\alpha \vee \beta$ is also not optimal. But $\alpha \vee \beta$ is implied by α , which, by hypothesis, *is* optimal.

A well-known example with this sort of structure is the notorious Professor Procrastinate (Jackson and Pargetter, 1986), who is asked to review a book but knows that,

³Some Consequentialists will reject the maximising aspect of this definition. They think an act can be permissible, though its consequences are not *best*, provided they are “good enough”. So far as I can tell, this view faces the same problem. The examples used to show that optimality is not closed under implication would equally show that having good enough consequences also lacks this property.

⁴An alternative would be to identify the consequences of an act with those of doing it in the *best* way possible, i.e., the way that would have best consequences. This would ensure that optimality is closed under implication. Suppose that act α implies act β . Then every way of doing α is a way of doing β . So the best way of doing β must have consequences at least as good as those of the best way of doing α . So if α is optimal then β is also optimal. I suspect, however, that this is really just a confusing way of stating Maximalism (see below). It achieves the same result, but only by adopting an idiosyncratic view of consequences.

were he to accept the review, he would not complete it by the deadline. His *accepting and doing the review* would have good consequences. Many would benefit: the book's author, the editor, potential readers, and so on. His *accepting the review*, however, would not have good consequences. Since he would not do the review if he accepted it, this would only delay the process, causing significant inconvenience. But his accepting and doing the review *implies* his accepting the review.⁵

In summary, the problem consists in a triad of jointly inconsistent claims:

- C1. An act is permissible if and only if it is optimal.
- C2. Permissibility is closed under implication.
- C3. Optimality is not closed under implication.

This has been called the Problem of Act Versions (so called because the different ways of performing an act can be called “versions” of it).⁶

Does the problem also occur for other moral theories? I shall not try to answer this question in any detail here, but *prima facie* it seems entirely possible that it does. On a Kantian theory, for example, permissibility is coextensive with what we might call “universalisability”, that is, the property of conforming to a maxim that can be rationally willed as a universal law. It is not obvious that universalisability must be closed under implication. It would at least take some work to show that it is. So Kantians may also have to contend with this problem. In what follows, however, I shall focus solely on Consequentialism. It is helpful to concentrate on a particular concrete example.

One solution to the Problem of Act Versions is to reject, or revise, C1. The trouble with this claim, one might think, is that it requires us to evaluate all acts in the same way, by their consequences. A better approach might be to evaluate only a privileged subclass of acts in this way, and to evaluate other acts by their relation to these privileged acts. Maximalism is such an approach. It privileges *maximal* acts, those that can be done in only one way, that have only one version.⁷ (In algebraic terms, it would be better to call these acts “atomic”, but I shall say “maximal” since this seems the standard terminology.) Thus Maximalism replaces C1 with:

C1*. An act is permissible if and only if it is implied by an optimal maximal act.⁸

For any property F , the property of being implied by an act that is F is closed under implication (this assumes only that implication is transitive). Thus C1* entails C2. This is how Maximalism solves the Problem of Act Versions. Still, Maximalism is controversial. Counterexamples have been proposed (I discuss two of these below). The question remains open, therefore, whether Maximalism is a satisfactory solution.

A general goal of this paper is to urge a somewhat novel approach to this question. Progress can be made here, I aim to show, by attending more carefully to the *structure* of acts (or more precisely: the structure of the act-implication relation). Many authors writing in this area seem to assume implicitly that acts have some sort of algebraic structure, and indeed some even employ formal algebraic notations.⁹ Few, however,

⁵A very similar example was earlier suggested by Goldman (1978). See also Carlson (1999, 253–254).

⁶Discovery of the problem is usually attributed to Bergström (1966) and Castaneda (1968).

⁷Maximalism was proposed by Goldman (1978).

⁸A maximal act is implied by only one maximal act, namely, itself. Thus C1* implies that a maximal act is permissible if and only if it is optimal.

⁹See, e.g., Gustafsson (2014).

explicitly discuss this structure. I want to suggest that they may be missing something important.

I shall employ considerations of structure to defend Maximalism in two ways. First, I shall propose a new argument for Maximalism. Insofar as we are concerned solely to solve the Problem of Act Versions, Maximalism is not compulsory. Other solutions are possible. For example, instead of privileging maximal acts, we might privilege “minimal” acts. Whereas a maximal act is one that is not implied by any other act, a minimal act is one that does not imply any other act.¹⁰ The view which might be called “Minimalism” holds that an act is permissible just in case it does not imply any sub-optimal minimal act. It is easy to see that Minimalism, like Maximalism, also entails that permissibility is closed under implication. Opting for Maximalism may therefore seem unmotivated. What’s so special about maximal acts? Why privilege them rather than minimal acts, or some other class of acts?

Moreover, it might seem that privileging *any* class of acts is inimical to the spirit of Consequentialism. This theory, as commonly understood, does not say *merely* that permissibility is coextensive with optimality. Rather, it says that the former is, in some sense, *explained* by the latter: an act is permissible *because* it is optimal. But on Maximalism (or, for that matter, Minimalism) this cannot in general be true. A non-maximal act may be permissible despite its failing to be optimal. The Maximalist seems therefore committed to a disunified account of permissibility: what explains the permissibility of maximal acts is not what explains the permissibility of non-maximal acts. If there is no rationale for such a disunified account apart from solving the Problem of Act Versions, then it may seem *ad hoc*. The view has simply been “rigged” to get the desired result.

What would help the Maximalist, then, is an independent argument for privileging maximal acts, that is, one that does not rest solely on solving the Problem of Act Versions. This is what I aim to provide. When we think about the structure of acts, we see that maximal acts really are special in an important way. Options are reducible to maximal acts.

Second, I shall respond to some objections to Maximalism. One objection is that maximal acts do not exist: every act is implied by another. Other objections involve putative counterexamples, wherein Maximalism seems to get the wrong result. I aim to show here, too, that thinking about the structure of acts can help to see where these objections go wrong.

As we shall see, there are issues about the structure of acts which are interesting in their own right, not only because of their bearing on the debate about Maximalism. Are there such things as maximal acts? Does any arbitrary collection of acts have a disjunction? Is it possible to perform an act without performing it in any specific way? The discussion of these issues below may be of interest even to readers who are not especially interested in Maximalism.

2 An Argument for Maximalism

The argument I shall propose rests on a formal result. This result has two main assumptions. The first concerns the structure of the act-implication relation. It says, in

¹⁰Strictly, this should say “any other *contingent* act”. An act is contingent if it is both possible to do the act and possible not to do it. An act is maximal if it is implied by no act except itself and \perp . An act is minimal if it implies no act except itself and \top . See below for discussion of the acts \top and \perp .

technical terms, that this relation has the structure of a “complete Boolean algebra”. The second assumption concerns the nature of an agent’s options. It says that these are complete and consistent sets of acts. As I shall show, these assumptions together imply that options are *reducible* to maximal acts. This is what makes maximal acts special. First, however, I shall explain each assumption in more detail.

2.1 The Structure of Act Implication

Roughly speaking, a Boolean algebra is the sort of structure we need in order to use the operations of conjunction (“and”), disjunction (“or”), and negation (“not”). A familiar structure of this sort is the “power set” of any set, i.e., the set of all subsets of this set, where these operations correspond to intersection, union, and complement. The premise, then, is that these operations can also be applied to *acts*. As well as the acts of *going to the pub* and of *drinking whisky*, for example, there are the acts of *going to the pub and drinking whisky*, and of *either going to the pub or not drinking whisky*, and so on.

More exactly, we have the following. Let \mathcal{A} be the set of all acts available to some agent at some time. And let \leq be a binary relation on \mathcal{A} , representing act implication; i.e., $\alpha \leq \beta$ means that α implies β . Then the assumption that (\mathcal{A}, \leq) is a complete Boolean algebra may be divided into four assumptions.

First, \leq is a partial order: reflexive, transitive, and antisymmetric. Only the last of these conditions could be contentious. Clearly, it is impossible to do α without doing α (reflexivity); and if it is impossible to do α without doing β and impossible to do β without doing γ , then it must be impossible to do α without doing γ (transitivity). Antisymmetry requires that numerically distinct acts cannot be practically equivalent, in the sense that it is impossible to do either without doing both. Perhaps there are counterexamples like the following. It is impossible for me to jump less than one metre without jumping less than two metres. And it is impossible for me to jump less than two metres without jumping less than one metre (because I cannot jump more than a metre). But jumping less than one metre is not the same as jumping less than two metres. From an ethical perspective, however, it seems harmless to treat these as the same. It is hard to see how practically equivalent acts could differ in their moral status.

Second, every set of acts $A \subseteq \mathcal{A}$ has a “join”, denoted $\bigvee A$, in (\mathcal{A}, \leq) . The join of A is an act such that (i) this act is implied by every act in A , and (ii) any other act that is implied by every act in A is implied by this act. In other words, it is the most specific act that is implied by all acts in A . We may think of $\bigvee A$ as the *disjunction* of A : the act of performing at least one of the acts in A .

Third, every set of acts $A \subseteq \mathcal{A}$ has a “meet”, denoted $\bigwedge A$, in (\mathcal{A}, \leq) . The meet of A is an act such that (i) this act implies every act in A , and (ii) any other act that implies every act in A implies this act.¹¹ In other words, it is the least specific act that implies all acts in A . We may think of $\bigwedge A$ as the *conjunction* of A : the act of performing all of the acts in A .¹²

It follows that the set of all acts, \mathcal{A} , itself has both a join and a meet. These are commonly called, respectively, “top” and “bottom”, denoted \top and \perp . In the present context, we can consider these the “necessary act” and the “impossible act”. Whatever

¹¹In technical terms, $\bigvee A$ is the supremum or least upper bound of A , and $\bigwedge A$ is the infimum or greatest lower bound of A .

¹²What makes the Boolean algebra *complete* is that joins and meets exist for *all* subsets of \mathcal{A} , not merely for finite subsets. A finite Boolean algebra is necessarily complete.

the agent does, she cannot do \perp , and cannot fail to do \top . Qualms about the existence of such acts will, I hope, be assuaged by noting that they are included here solely for technical convenience. We could do without them, but at the cost of a less elegant framework.¹³

Fourth, every act $\alpha \in \mathcal{A}$ has a “complement”, denoted $\neg\alpha$, in (\mathcal{A}, \leq) . The complement of α is an act such that (i) the disjunction of α and this act is \top , and (ii) the conjunction of α and this act is \perp . The agent must do either an act or its complement, but cannot do both. We may think of $\neg\alpha$ as the “negation” of α : the act of not doing α .

2.2 Heterogeneous Acts

Why think acts have such a structure? One reason is that Boolean operations seem quite common and natural in ordinary discourse about acts. We seem happy to talk about “going to bed *and* reading a book”, and about “*not* staying up late”, and so on. This suggests that our common way of thinking about acts presupposes something like this structure.

Still, it might be denied that constructions of this kind correspond to composite acts. It might be claimed, for example, that “not staying up late” refers not to an act, but to the absence of one; or that “going to bed and reading a book” refers to two acts, not one. Most controversial, I expect, will be disjunctions. The sceptic may accept that *raising your hand* is an act, and so is *stamping your foot*, but deny that there is any such act as *either raising your hand or stamping your foot*, for example. This description may seem too heterogeneous to constitute a single act. Suppose that you raise your hand (without stamping your foot), while I stamp my foot (without raising my hand). It might seem absurd to say that you and I have thereby performed the same act.

However, the apparent absurdity in saying this can be explained in a way that is consistent with the existence of such an act. It would be odd to say that Kant and Hume were born in the same place. Kant was born in Königsberg, Hume in Edinburgh. These are different places. Still, Kant and Hume *were* both born in Europe. This should not lead us to conclude that there is no such place as Europe. There is no real mystery here. Actually, Kant was born in many places. He was born in Königsberg, and in Germany, and in Europe, and so on. In conversation, context selects one of these places, at an appropriate level of specificity, to be considered “the” place of Kant’s birth. Typically, perhaps, we expect a person’s place of birth to be given as a city, not a country or continent. But it is not hard to imagine contexts requiring either greater or lesser specificity. In some cases, for example, we would not count two people as being born in the same place unless they were born in the same hospital. Likewise, we may say, though there is such an act as *either raising your hand or stamping your foot*, this act is at too low a level of specificity to be considered “the” act that either you or I performed. This is why it seems odd to say that we performed the same act.

Notice, furthermore, that disjunctive acts need not be so artificial. Strictly speaking, *every* non-maximal act is disjunctive. If an act can be done in several ways, then, by definition, it is the disjunction of these ways. Thus even *raising your hand* — which the sceptic, I assume, is happy to admit as an act — is disjunctive. It is the disjunction of *raising your hand slowly*, *raising your hand quickly*, and so on. A blanket ban on

¹³For this reason, it is not strictly correct to say that \mathcal{A} is the set of acts available to the agent, since the impossible act is not in any normal sense “available”. We should rather say that \mathcal{A} contains all the available acts, plus the impossible act.

disjunctive acts would therefore have the effect of eliminating all but maximal acts. This would be a more severe form of Maximalism.

An intermediate view allows disjunctive acts, but only in special cases: some collections of acts form disjunctions while others do not. But what could make the difference? Perhaps acts form a disjunction only when they are suitably unified, when they are all different ways of doing some common thing. The acts *raising your hand slowly* and *raising your hand quickly* pass this test, since both are ways of *raising your hand*. This argument, however, seems to presuppose what it aims to establish. Were there such an act as *either raising your hand or stamping your foot*, then *raising your hand* and *stamping your foot* would each be ways of doing this.

We can, moreover, give a positive argument for disjunctive acts. Our interest in acts, in this context, is a moral one. Acts are the objects of a certain sort of moral evaluation: roughly, evaluation of what people *do*.¹⁴ These objects may often be disjunctive in nature. We may consider, for instance, whether I am permitted to stay home from work today. This question covers a broad array of things that I might do if I stayed home: I could play video games, or clean the windows in my flat, or write an overdue book review, or catch up on my email correspondence, or It makes sense, then, to think of this question as calling for moral evaluation of a disjunctive act, the disjunction of all these things I might do. There may even be circumstances in which we are concerned to know whether a person is permitted to either raise her hand or stamp or foot; so we should count this as an act too.

Suppose, for instance, that two cyclists are riding very quickly towards a blind intersection. You can see that, unbeknown to either cyclist, they are on a collision course. Your only hope to avert disaster is to signal to one of them to stop (you haven't enough time to signal to both). You have established previously that the stop signal for one cyclist is raising your hand, and for the other is stamping your foot. It makes no difference which cyclist stops; the consequences will be no better or worse either way. So you ought to either raise your hand or stamp your foot. You have an obligation to this. But what is the content of this obligation? It cannot be the act of raising your hand, because it's okay for you not to do this, if you stamp your foot instead. Likewise, it cannot be the act of stamping your foot. If the content of your obligation is an act, then it must be the act of either raising your hand or stamping your foot.¹⁵ If it is not an act, then it is nonetheless the sort of thing to which a moral theory like Maximalism can be applied. Obviously, a similar example could be constructed for any arbitrary collection of acts. Therefore, we have reason to include arbitrary disjunctions of acts among the objects of moral assessment. It matters little whether we call these disjunctions "acts". I would be happy to call them something else if I knew of a better name.

2.3 Options

I turn now to the second assumption, regarding the nature of options. We are accustomed to thinking of options as *single* acts. The agent faces a menu of mutually exclusive acts from which she must choose exactly one. Such a picture seems inadequate, however, if acts are structured as described above. Whatever the agent does, she will do not just one act, but many. She will not merely raise her left arm, but will also raise one of her

¹⁴There may be other kinds of moral evaluation, for example, evaluation of how people *are*, whose objects are something like character traits or dispositions.

¹⁵Thanks to an anonymous reviewer for suggesting this argument.

limbs, and move her body, and signal to the taxi driver, and so on.¹⁶ Acts come grouped together as “package deals”. The items on the menu are, therefore, better understood not as single acts, but as collections or sets of acts.

This might seem not to reflect the phenomenology of choosing. Even if an agent will inevitably *do* many acts, she might not *choose* to do many acts. Choosing seems seldom to involve consciously identifying a plurality of acts to perform. I may declare my choice, either to myself or someone else, by simply saying: “I’m going to watch the latest episode of *Veep*.” I need not say: “I’m going to simultaneously do each of the following: watch the latest episode of *Veep*; watch a TV show in the political satire genre; either watch TV or read a novel; engage in leisure activity; . . .” But we must distinguish between choosing to do an act and declaring that one will do it. The former does not require the latter. I know that *Veep* is a political satire. Thus my watching a political satire is no less a matter of choice than my watching *Veep*. It’s not as if my watching *Veep* is voluntarily, while my watching a political satire happens against my will or as some unforeseen side-effect; it is part of what I choose to do, though I might not consciously and explicitly identify it as such.¹⁷

So options are sets of acts. But not every set of acts is an option. Whatever the agent does, there will be a set of acts such that she does all and only the acts in this set. An option, we may say, is a candidate for being this set. It is a set of acts such that it is possible to do all and only the acts in this set. When it is possible to do *all* the acts in a set, we may say that the set is *consistent*. When it is possible to do *only* the acts in a set, we may say that the set is *complete*. The premise, then, is that an option is a consistent and complete set of acts.

Formally, these conditions may be defined as follows. First, a set of acts A is consistent if and only if $\bigwedge A \neq \perp$. As I said above, to do all the acts in A is to do $\bigwedge A$. Thus it is possible to do all the acts in A just in case doing $\bigwedge A$ is not impossible, i.e., $\bigwedge A \neq \perp$. Second, A is complete if and only if $\bigvee \bar{A} \neq \top$. \bar{A} is the set of all acts not in A . To do only the acts in A is to not do any act in \bar{A} , and this is to not do $\bigvee \bar{A}$. Thus it is possible to do only the acts in A , just in case $\bigvee \bar{A}$ is not necessary, i.e., $\bigvee \bar{A} \neq \top$.

2.4 Intentional Action vs Mere Behaviour

It seems hard to deny that options must be both consistent and complete. Surely an option must include *all* the acts the agent will perform if she chooses this option, and *only* these acts. However, it might be objected that the formal definitions of consistency and completeness I have assumed do not adequately capture these notions. One might think, for example, that a set of acts which is incomplete, according to my definition, may nonetheless be such that it is possible to do only the acts in this set.

It may be helpful to note that a set of acts A is both consistent and complete (in the sense I have defined) just in case it satisfies the following four conditions:

1. $\perp \notin A$
2. $\top \in A$

¹⁶By “acts” here I mean act *types*. When I say “the agent will do many acts”, I mean, more precisely, that there are many types of act such that the agent will do an act of this type.

¹⁷Of course, it is possible to choose to watch *Veep* without knowing that it is a political satire. But such cases, where the agent lacks full information, are not relevant here. A background assumption in the Maximalism debate is that we are concerned with “objective” permissibility.

3. If $B \subseteq A$ then $\bigwedge B \in A$

4. If $B \subseteq \bar{A}$ then $\bigvee B \notin A$

The first and second conditions seem uncontroversial. As I've said, whatever the agent does, she cannot do \perp (the impossible act) and cannot fail to do \top (the necessary act); so the latter must be included, and former must not be included, in all options. The third seems okay too. If the agent does all the acts in some set then she does the conjunction of these acts. However, for reasons I shall explain below, the fourth condition may seem doubtful in light of the familiar distinction between intentional action and mere behaviour.

The condition says that option-membership “distributes over” disjunction. In the finite case, it is equivalent to the following: $\alpha \vee \beta \in A$ only if either $\alpha \in A$ or $\beta \in A$. The underlying idea, then, is that if an act can be done in more than one specific way, then it cannot be done without doing it in one of these specific ways. For example, you cannot drink whisky without either drinking Lagavulin, or drinking Laphroaig, or If you do not drink any particular kind of whisky, then you do not drink whisky at all. (Compare: an object cannot simply be red, without being some particular shade of red, i.e., scarlet, or crimson, or)

To see why this might be problematic, consider the example of throwing a dart at a dartboard. To hit the dartboard, you must hit it in some particular region: you must either hit the bullseye, or hit the double twenty, or In some cases, however, a region may be too small (or too dispersed) for your hitting it to be counted as an act you have performed. Suppose, for instance, your aim is not especially good: if you try to hit the bullseye, you will succeed, on average, only once in a hundred attempts. But on one lucky occasion, by pure fluke, you happen to hit the bullseye. Although, in one sense, hitting the bullseye is something you did, it was not sufficiently within your control to be counted as an act that you performed. It is not appropriately attributable to *you*. You don't deserve praise for hitting the bullseye (nor would you have deserved blame for missing it, provided you had at least tried your best to hit it). In terms of the common distinction, hitting the bullseye was a *mere behaviour*, not an *intentional action*.

Now this might suggest a counterexample to the condition in question. Suppose the following is true: when you aim at a given point on the dartboard, the dart reliably hits somewhere within a one-inch radius of that point. Thus, you could hit the left half of the dartboard, if you chose to do so, by aiming at a point in the middle of the left half. That is to say, hitting the left half of the dartboard would, in this case, count as an act you performed. Likewise, you could hit the right half of the dartboard. But now suppose that you aim for the dead centre of the dartboard. Then whichever side the dart lands on, left or right, you do not perform the act of hitting the left half, nor the act of hitting the right half, because you do not exercise sufficient control to ensure either of these outcomes. But you *do* perform the act of hitting either the left half or the right half of dartboard, since this is just equivalent to hitting the dartboard, and you do ensure this outcome.¹⁸ This looks like a counterexample: you perform the disjunction of two acts without performing either act.

As I shall argue, however, this is not really a counterexample. Very roughly, we need to distinguish two acts: one is the act of *hitting either the left half or the right half*; the other is the act of *either hitting the left half or hitting the right half*. Whereas the former

¹⁸I assume here that the vertical diameter belongs both to the left half and the right half.

is performed by you, the latter is the disjunction of the acts you refrain from performing. The apparent counterexample depends on conflating these two acts.

To make this more precise, it will be helpful to introduce the notion of “targeting” a region of the dartboard. To target a region X , in this sense, is to throw the dart in such a way that it cannot hit any point outside X , but can hit any point inside X .¹⁹ It is to exercise *just enough* control to ensure that the dart hits somewhere within the region X , and no more: the dart must hit *some* point within X , but which particular point it hits is left to chance (or to causal factors beyond the agent’s action). Let $t(X)$ be the act of targeting X .

To *hit* a region, we may then say, is to target some region within this region. That is, if $h(X)$ is the act of hitting X , then $h(X) = \bigvee\{t(Y) : Y \subseteq X\}$. Suppose you target X and the dart hits Y . Then there are two cases. If $X \subseteq Y$ then, because you ensure that the dart hits X , you also ensure that it hits Y . So hitting Y is attributable to you, an act you perform. On the other hand, if $X \not\subseteq Y$ then, as far as your action goes, it is merely a matter of luck that the dart hits Y . So, in this case, hitting Y is not an act you performed.

Now, let L and R be, respectively, the left half and right half of the dartboard. As the above example shows, it is possible for you to hit the disjunction (or union) of these regions without hitting either region. That is, you can perform $h(L \cup R)$ without performing either $h(L)$ or $h(R)$. You do this by targeting some region contained within $L \cup R$, but not within either L or R , e.g., the top half of the dartboard. But the act of hitting the disjunction of two regions is not the same as the disjunction of the acts of hitting the two regions: $h(L \cup R) \neq h(L) \vee h(R)$. We have: $h(L \cup R) = \bigvee\{t(X) : X \subseteq L \cup R\}$ and $h(L) \vee h(R) = \bigvee\{t(X) : X \subseteq L \text{ or } X \subseteq R\}$. These are not the same, precisely because some regions contained within $L \cup R$ are not contained within either L or R . If you target such a region, you will perform $h(L \cup R)$, but not $h(L) \vee h(R)$. Moreover, it is clear that you cannot perform $h(L) \vee h(R)$ without performing $h(L)$ or $h(R)$.

So it seems that the condition in question is acceptable even when Maximalism is restricted to intentional acts, as of course it should be.

2.5 Options are Reducible to Maximal Acts

These assumptions entail that options are reducible to maximal acts, in the following sense. For any act α , let $\uparrow\alpha$ be the set of all acts implied by α , i.e., $\uparrow\alpha = \{\beta : \alpha \leq \beta\}$. In technical terms, $\uparrow\alpha$ is a *principal filter*, and α is its *principal element*. Then options are just principal filters whose principal elements are maximal. That is, A is an option if and only if there exists a maximal act α such that $A = \uparrow\alpha$ (see Appendix for proof). Thus there is a perfect one-to-one correspondence between options and maximal acts. Every option contains exactly one maximal act, and every maximal act is contained in exactly one option.

This explains what is special about maximal acts, why it is not *ad hoc* to privilege them in the way that Maximalism does. The same is not true, for example, in the case of minimal acts. Every option contains many minimal acts, and every minimal act is contained in many options. To decide on a maximal act is to fully determine what to do (at least insofar what you do is within your control). It is to settle, for each act, whether or not to do it. But this is not so for other acts, and especially not for minimal acts. After deciding only on a non-maximal act, you still have some deciding to do, since

¹⁹Clearly this is an idealisation. A more realistic model would employ a probability distribution over regions of the dartboard. But the simpler model is sufficient for present purposes.

you have yet to decide in what way to do this act. In effect, the agent's choice between options is a choice between maximal acts.

In order to establish Maximalism, two further premises are needed. First, an option A is permissible if and only if $\bigwedge A$ is optimal (recall that $\bigwedge A$ is the act of doing every act in A).²⁰ Second, an act is permissible if and only if it is contained in at least one permissible option.²¹ Putting these two premises together, we have the following: an act α is permissible if and only if, for some option A , $\bigwedge A$ is optimal and $\alpha \in A$. If options are reducible to maximal acts, as shown by our formal result, this is equivalent to the following: an act α is permissible if and only if, for some maximal act β , $\bigwedge \uparrow\beta$ is optimal and $\alpha \in \uparrow\beta$. But we know that $\bigwedge \uparrow\beta = \beta$. So, we have Maximalism: an act is permissible if and only if it is implied by an optimal maximal act.

I have little to say here in defence of these premises, except that they seem plausible to me. (Every argument must start somewhere.) The point I wish to emphasise, however, is the crucial role played in the argument by considerations of the formal structure of acts. Without the formal result, these further premises would be insufficient to establish Maximalism. This helps to show, I hope, some of the utility of thinking about structure in these matters.

3 Objections

Thinking about structure may also be beneficial in another way: it may help the Maximalist respond to objections. In this section I shall consider three objections. The first is simply that there are no such things as maximal acts. The second and third both involve purported counterexamples, one advanced by Bykvist, the other by Johan Gustafsson. In each case, I aim to show that careful consideration of the structure of acts reveals what is wrong with these objections.

3.1 Are There Maximal Acts?

Clearly Maximalism depends on the existence of maximal acts. Without them, there would be no options; every set of acts would be either inconsistent or incomplete (see Appendix). Maximalism would then imply, absurdly, that *no* act is permissible. But some may be sceptical of the notion of a maximal act. It may seem that, no matter how much detail we include in our specification of an act, it will always be possible to specify it even more.

Kristen Bykvist considers, and rejects, an argument against the existence of maximal acts:

Surely any adequate act-ontology must count actions as well as inactions (omissions). It seems natural to suppose that we can form act-compounds

²⁰Were we beginning with a moral theory other than Consequentialism, this premise could be adjusted accordingly. For example, were we developing a Maximalist form of Kantianism, the premise would be: an option A is permissible if and only if $\bigwedge A$ is universalisable. The conclusion would then be that an act is permissible if and only if it is implied by a universalisable maximal act.

²¹It may be tempting to think that, for an act to be permissible, *every* option that contains it must be permissible. For example, if the bartender says "You may order whisky", we would normally take this to imply that you may order *any* kind of whisky (available in the pub), not merely that there is at least one kind you may order. But this can be explained pragmatically, in terms of conversational implicature. If you are permitted to order only some kinds of whisky, then we expect the bartender to say something more informative, e.g., "You may order any whisky except Laphroaig."

of actions and inactions. But then we can construct the following list of act-compounds:

- S 's raising his hand at t and not doing a_1
- S 's raising his hand at t and not doing a_1 and not doing a_2
- S 's raising his hand at t and not doing a_1 and not doing a_2 and not doing a_3
- \vdots

Since whatever I do I will always fail to perform an infinite number of actions this list will contain an infinite number of distinct actions, each of which is entailed by the subsequent ones. Hence we will not have any maximal actions. (Bykvist, 2002, 64)

Bykvist responds by denying one of the assumptions of the argument (that there can be infinitely many acts which the agent intends not to perform). But this is unnecessary, because the argument is in fact invalid. The conclusion, that there are no maximal acts, simply does not follow. All that follows is that no act *in the list* is maximal. But this does mean that no act *at all* is maximal, unless we assume that every act must appear somewhere in the list. I see no reason to accept this further assumption, and Bykvist does not offer one.

The basic point here is that infinite divisibility is compatible with the existence of atomic elements. Consider, for example, the real interval $[0, 1]$.²² This may be divided into two sub-intervals, $[0, \frac{1}{2}]$ and $[\frac{1}{2}, 1]$, each of which may then be divided into two further sub-intervals, and so on forever. So we have here infinite divisibility. Nonetheless, this interval does have atomic parts, namely, singleton sets such as $\{0\}$. That we can go on dividing forever, without ever reaching anything indivisible, does not imply that nothing is indivisible. Similarly, as Bykvist's example shows, acts may also be infinitely divisible. We may be able to construct infinite sequences of increasingly specific acts. But this does not imply that there are no maximal acts.

Still, the idea of a maximal act might not seem very intuitive. We find ourselves hard pressed to give an example of one. Descriptions of acts can be *refined*. For example, "drinking whisky" can be refined to "drinking whisky while sitting", which can be refined again to "drinking whisky while sitting in an armchair", and yet again to "drinking whisky while sitting in an armchair with tartan upholstery", and so on. Each successive refinement describes an act that implies but is not implied by its predecessors. This process seems like it could go on forever. Moreover, it seems that *any* description of an act must be refinable. A description so detailed that it resists any further specification is hard to imagine.

A refinable description cannot describe a maximal act. So, if every description is refinable, then there are no *describable* maximal acts. But it does not follow that there are no maximal acts at all. For there could yet be some indescribable ones. Our inability to give examples of maximal acts may be due, not to their absence, but rather to limitations in the descriptive power of our language. Here is a rough analogy. It is impossible for us to write down the decimal expansion of π , because π is infinite and we finite. We can write down each of a sequence of decimal expansions — 3, 3.1, 3.14, . . . — each one a more accurate approximation of π than the previous. But we will never

²² $[x, y]$ is the set of every real number no less than x and no greater than y , i.e., $[x, y] = \{z : x \leq z \leq y\}$.

reach one that cannot be further refined by appending another digit at the end. This surely is no reason to doubt that π exists (we may have other reasons). What there is and what we can describe are separate questions.

This defence of maximal acts may seem unsatisfactory in the present context. According to Maximalism, maximal acts are the fundamental bearers of moral properties such as permissibility. If these are indescribable, then it may be hard to see how moral discourse or thought is even possible. Are we really talking fundamentally about things that we cannot even describe? But here it is useful to recall the distinction introduced earlier between acts and the more general class which we might call “doings”. Only those of our doings which are sufficiently within our control count as acts. Even when you do something as simple as move your arm, the very precise physical movements your body undergoes may be too complex to be perfectly accurately described in a finite sentence of English. However, such precise movements are not within your control anyway. So they are not acts. It is perhaps more plausible that the limits of what is within our control will be reached before the limits of what we can describe.²³

These arguments against maximal acts are unpersuasive. Let me now suggest an argument in favour maximal acts. Suppose a person acts on some occasion. Let X contain *all* and *only* the acts she performs. Now, since actuality implies possibility, it must be *possible* to perform all and only the acts in X . But that is just to say that X is consistent and complete; that it is an option, in our terms. It follows that $\bigwedge X$ is maximal (see Appendix, Lemma 2). So, on the assumption that the agent acts, there must be at least one maximal act. Without maximal acts, action would be impossible.

In any case, there may be a way of revising Maximalism so that it does not require maximal acts. Here’s the basic idea. Even if it were possible to continue refining act descriptions forever, we would expect eventually to reach a point after which further refinement could make no *moral* difference. For example, it makes no moral difference whether the murderer shoots her victim while raising her left eyebrow or her right. Since the consequences are the same either way, these acts are morally equivalent. The rough idea, then, is that acts described at this level of refinement, though they may not strictly be maximal, may nonetheless be treated as such for moral purposes.

This can be made more precise as follows. First, say that an act α is *morally basic* if every act that implies α is morally equivalent to α (i.e., has the same consequences as α , or at least has consequences that are no better or worse than those of α). Plausibly, non-maximal acts may be morally basic. That is, there may be an act such that, though it can be done in several ways, it doesn’t matter, morally, in which way it is done. Now say that a partition of the space of acts, $\alpha_1, \alpha_2, \dots$, is *morally neutral* if every α_i is morally basic. We may then say that α is permissible according to a morally neutral partition if it is permissible in the sub-algebra that would result from discarding all acts that are not implied by at least one α_i . Finally, we may say that an act is permissible if it is permissible according to some morally neutral partition. This version of Maximalism does not require maximal acts. It requires only morally basic acts, maximal or not.

3.2 Bykvist’s Counterexample

Bykvist (2002, 55) presents a counterexample to Maximalism, shown in Figure 1. This represents various acts or sequences of acts that an agent may perform in a particular situation, along with the consequences of his performing them. The thicker line indicates that “the agent would freely do a_3 if he were to do a_1 ” (Bykvist, 2002, 50). The dashed

²³This argument is made by (Portmore, 2015).




Figure 1: Bykvist's example

line indicates that $a_1 \wedge a_3$ is not available to the agent at t_1 , a_3 becoming available only at t_2 . The consequences of these acts are as shown: if the agent does a_3 , there will be no disaster; if he does a_4 , there will be a major disaster; and if he does a_2 , there will be a minor disaster.

The question is what the agent should do at t_1 . Bykvist argues that Maximalism gives an absurd answer:

Here the maximal actions are $[a_1 \wedge a_4]$ and a_2 , and both will have disastrous consequences. (Recall that $[a_1 \wedge a_3]$ is not performable.) [Maximalism] would therefore claim that a_2 is obligatory since, of all the maximal actions, this action would have the best consequences (it will only lead to a minor disaster). But this prescription is absurd since there is a way to avoid a disaster, namely, to perform the non-maximal action a_1 . (Bykvist, 2002, 55).

But there is something odd in Bykvist's description of the example. To see what it is, it will be helpful first to consider an analogy from mereology, the theory of "parthood" relations. In mereology, it is common to treat parthood as a reflexive relation: everything is a part of itself. The perhaps more familiar, irreflexive notion of parthood is called "proper parthood", defined as follows: x is a proper part of y if and only if x is a part of y , and y is not a part of x . Now, a very plausible and widely accepted mereological principle is that nothing can have only one proper part.²⁴ There may be things with zero proper parts, namely, atoms. But if a thing has *any* proper parts then it must have *many*, at least more than one. Thus, a tree has many branches, a year has many days, a book has many chapters, and so on. We can, of course, imagine a book with only one chapter (and no prefaces, forewords, appendices, etc.). But then that chapter would not be a *proper* part of the book; it would be the whole thing.

The implication relation between acts is much like parthood. We can define a notion of "proper versionhood", analogous to proper parthood, as follows: α is a proper version of β if and only if α is a version of β , and β is not a version of α . (Recall that α is a version of β if α implies β , i.e., it is impossible to do α without doing β .) An analogous principle would then say that no act can have only one proper version. If an act has any proper versions (i.e., it is non-maximal) then it must have many. For example, visiting Germany has many proper versions: visiting Munich, visiting Berlin, and so

²⁴See e.g. (Simons, 1987, 26-8).

on. Likewise, drinking whisky has many proper versions: drinking Lagavulin, drinking Laphroaig, and so on. We can, of course, imagine a situation in which, say, the agent is not able to visit any German city except Munich. But in this situation, visiting Munich would not be a *proper* version of visiting Germany. These acts would be practically equivalent: it would be impossible to do one without also doing the other.

This principle, like the analogous mereological principle, seems very plausible. However, it is violated by Bykvist's example, as he describes it: the act a_1 has only one proper version, $a_1 \wedge a_4$. None of the other acts mentioned by Bykvist is a proper version of a_1 . This problem would be solved by including the act $a_1 \wedge \neg a_4$, since it would be another proper version of a_1 . But this would also undermine Bykvist's argument. He must assume that $a_1 \wedge a_4$ and a_2 are the only maximal acts; otherwise he could not infer that a_2 has the best consequences from fact that a_2 has better consequences than $a_1 \wedge a_4$. But were $a_1 \wedge \neg a_4$ included, it, too, would be maximal. Moreover, it would have better consequences than a_2 : were the agent to do $a_1 \wedge \neg a_4$, he would go on to do a_3 , thereby averting disaster. Thus Bykvist could no longer saddle the Maximalist with the implausible conclusion that a_2 is obligatory. Maximalism would imply that $a_1 \wedge \neg a_4$, and therefore a_1 , is obligatory — the intuitively correct conclusion.²⁵

It might seem inconsistent to say that the agent can do $a_1 \wedge \neg a_4$ at t_1 . This entails that the agent can do $\neg a_4$ at t_1 , which might seem to foreclose his doing a_4 at t_2 ; but, by hypothesis, his doing a_4 must remain a possibility at t_2 , unless he does a_2 at t_1 . However, while doing both $\neg\alpha$ and α at the *same* time is indeed impossible, doing $\neg\alpha$ at one time need not preclude doing α later. To do $\neg\alpha$, as I have said, is just to not do α . But, of course, one can often do tomorrow what one does not do today (otherwise we might not be tempted to procrastinate so much). Thus the agent's doing $a_1 \wedge \neg a_4$ at t_1 is compatible with his doing a_4 at t_2 .

An analogy may be helpful. Suppose there is an Irn-Bru vending machine before you.²⁶ The machine dispenses cans of Irn-Bru at \$0.50 each, and you have \$1.00 to spend; so you can buy either one or two cans. Now, we might describe your options as follows. Let c_1 be the act of buying the first can (i.e., the next one to be dispensed), and let c_2 be the act of buying the second can. Then you can either do c_1 or do $c_1 \wedge c_2$. But then c_1 is implied by $c_1 \wedge c_2$ and not by any other maximal act. Thus it seems that c_1 has only one proper version.

The mistake here, however, is identifying the act of *buying one can* with that of *buying the first can*.²⁷ These are not the same. Although buying one can implies buying the first can, the converse implication does not hold. Buying the first can is compatible with also buying the second can, but in that case, you buy *two* cans, not one. Thus it is possible to buy the first can without buying one can. To buy one can is, rather, to buy *only* the first can, to buy the first *and not the second*, or $c_1 \wedge \neg c_2$. When this act is added, c_1 has two proper versions. Notice also, if you choose to buy only the first can now, this does not prevent you from returning later to buy the second. That is, your doing $c_1 \wedge \neg c_2$ now is compatible with your doing c_2 later.

²⁵Gustafsson (2014, 588, n. 9) makes a similar objection.

²⁶Irn-Bru is a Scottish brand of carbonated soft drink, or "juice", as such drinks are known in Scotland.

²⁷Perhaps "buying one can" is ambiguous between buying *exactly* one can and buying *at least* one can. Buying the first can *is* equivalent to buying at least one can. But the relevant option is buying exactly one can.

3.3 Gustafsson's Objection

Johan Gustafsson's objection to Maximalism is based on the following example:

Suppose, for example, that the agent gets an offer to participate in a version of Newcomb's problem. In this problem, there are two boxes. The first contains \$1,000, and the second contains either \$1,000,000 or nothing. The agent is given a choice between either taking what is in both boxes or only taking what is in the second box. The contents of the second box depends on what a predictor predicts the agent will choose. If the predictor predicts that the agent will take what is in both boxes, he puts nothing in the second box. If the predictor predicts that the agent will take only what is in the second box, he puts \$1,000,000 in the second box. . . . In this version of Newcomb's problem, the predictor bases his prediction on an observation of the agent at [the time when the agent either accepts or declines the offer]. By either a brain scan of the agent . . . or a meticulous observation of the agent's behaviour . . . , the predictor can tell whether the agent has formed at [that time] an intention to [accept the offer and take what is in both boxes] or if the agent has merely formed at that time an intention to [accept the offer]. (Gustafsson, 2014, 592)

The problem for Maximalism, says Gustafsson, is that it cannot differentiate between two possibilities:

- P1. The agent accepts the offer and then later takes what is in both boxes, but does not form the intention to take what is in both boxes until after she accepts.
- P2. As above, except the agent forms the intention to take what is in both boxes at or before the time when she accepts.

Since the prediction is based on the intentions of the agent at the time when she accepts, the agent receives \$1,001,000 in P1, but only \$1,000 in P2. Thus, what the agent does in P1 has better consequences than what she does in P2 (we may assume that the predictor would otherwise donate the \$1,000,000 to some evil cause).

By saying that Maximalism cannot "differentiate" between these possibilities, Gustafsson means that it cannot prescribe one without also prescribing the other (Gustafsson, 2014, 592). This is a problem, he claims, because a plausible form of Consequentialism should be able to prescribe only P1. Gustafsson's argument, as I understand it, is the following. Maximalism prescribes a possibility if and only if the agent performs some permissible maximal act in this possibility. But there is no maximal act that the agent performs in P1 but not in P2. Therefore, if Maximalism prescribes P1 then it also prescribes P2.

I believe this objection rests on an overly narrow understanding of Maximalism. The crucial issue is whether the agent acts differently in the two scenarios. Clearly, what the agent does in P1 is not entirely the same as what she does in P2. But is this a difference in the *acts* she performs? As we noted earlier, not everything you do is an act you perform.

According to Gustafsson, the difference is "the time at which the agent decides to [take both boxes]" (Gustafsson, 2014, 591). So it is a difference in the agent's decisions. A common view — perhaps the standard view — in the philosophy of action, is that

“practical decisions are momentary intentional actions of intention formation” (Shepherd, 2015, 336). On this view, there *is* a difference in the acts performed by the agent in the two scenarios, specifically at the time when she accepts: in P2, she performs, at this time, the act of deciding to take both boxes, whereas in P1 she does not. But in that case, Maximalism *can* differentiate the two possibilities. As I argued above, when a person acts, there must be some maximal act such that the agent performs all and only the acts implied by this act. It follows that, if there is any difference in the acts performed by the agent in P1 and P2, then there must be a difference in the *maximal* acts she performs. And that is all Maximalism needs to differentiate the two scenarios.

Gustafsson might reject the view that decisions are acts. Acts are things that you decide to do, he might say, but you do not decide to decide; you just decide. If that is so, however, then I see no reason why Maximalism should be restricted to acts. Whether they are acts or not, decisions surely are among the things people do which are properly subjected to moral scrutiny. Decisions may be right or wrong. (Were Gustafsson to deny this, it would be hard for him to maintain that there is any moral difference between the two possibilities, since, according to him, these differ only in the decisions of the agent.) Therefore, decisions may be included within the scope of a moral view like Maximalism.²⁸

It is worth noting that the alternative view proposed by Gustafsson, “Combinative Act Consequentialism”, seems very similar to Maximalism. He defines this view as follows:

It is obligatory for P in S to jointly intentionally perform all and only the acts in a set of acts X if and only if

1. P in S can jointly intentionally perform all and only the acts in X and
2. for every other set of acts Y such that P in S can jointly intentionally perform all and only the acts in Y , the outcome of P in S jointly intentionally performing all and only the acts in X is better than the outcome of P in S jointly intentionally performing all and only the acts in Y . (Gustafsson, 2014, 593)

Notice this refers to a set of acts such that it is possible to “jointly intentionally” perform all and only the acts in this set. This is essentially my definition of an *option*. Paraphrasing a little, Gustafsson’s view can, therefore, be stated more simply as follows: an act is obligatory if and only if it is contained in a option such that the consequences of performing this option are better than those of performing any other option. And this is perfectly compatible with Maximalism.

4 Conclusion

I have argued that Maximalism provides a solution to the Problem of Act Versions. Adopting this view enables us to give an account of cases like that of Professor Procrastinate while respecting the plausible principle that permissibility is closed under implication. However, this account may seem unsatisfactory for other reasons. According to Maximalism, Prof. P ought to do the review and therefore he ought to accept it (because he cannot do it without accepting it). But many people judge intuitively that he ought *not* to accept. I believe there is more to be said in defence of Maximalism

²⁸Portmore (2015) makes a similar response to Gustafsson’s objection.

here: it can, to a large extent, be squared with this common intuition. But that argument must wait for another time, since it would take us too far afield from the principal focus of this paper. The defence of Maximalism I have presented here is to this extent inconclusive. Still, I hope to have at least made a case for reorienting the debate so that questions about the structure of acts are given a more prominent role.

Appendix

The purpose of this appendix is to prove the main result advertised above, i.e., that in a complete Boolean algebra of acts, options are reducible to maximal acts, in the way specified.

We begin with some definitions.

Definition 1 (Acts). Let (\mathcal{A}, \leq) be a complete Boolean algebra.

Definition 2 (Maximal Acts). α is *maximal* if $\alpha \neq \perp$ and for all β , if $\beta < \alpha$ then $\beta = \perp$.

Definition 3 (Options). A is an option if $\bigwedge A \neq \perp$, and $\bigvee \bar{A} \neq \top$

Algebraic operations are defined in the standard way, e.g., for $A \subseteq \mathcal{A}$, $\bigvee A$ is the supremum of A in (\mathcal{A}, \leq) . Below $\uparrow\alpha$ denotes the principle filter whose principal element is α .

We may now prove the following:

Proposition 1. *A is an option if and only if there exists a maximal act α such that $A = \uparrow\alpha$.*

Proof. To prove the left-to-right implication, we show that if A is an option, then each of the following hold:

1. $\bigwedge A$ is maximal.
2. if $\alpha \in A$ then $\bigwedge A \leq \alpha$, i.e., $A \subseteq \uparrow\bigwedge A$.
3. if $\bigwedge A \leq \alpha$ then $\alpha \in A$, i.e., $\uparrow\bigwedge A \subseteq A$.

To prove (1), suppose $\alpha < \bigwedge A$. We need to show that $\alpha = \perp$. Since $\bigwedge A$ is a lower bound of A , and \leq is antisymmetric, it follows that $\alpha \notin A$. So, because A is complete, $\neg\alpha \in A$. But then $\alpha \leq \bigwedge A \leq \neg\alpha$. So $\alpha = \perp$.

To prove (2), it suffices to note that, by the definition, $\bigwedge A$ is a lower bound of A .

To prove (3), assume, for reductio, that $\bigwedge A \leq \alpha$ and $\alpha \notin A$. It follows that $\neg\alpha \in A$, otherwise A would be incomplete. But then $\bigwedge A \leq \alpha$ and $\bigwedge A \leq \neg\alpha$. So $\bigwedge A = \perp$, i.e., A is inconsistent.

To prove the right-to-left implication, we need to show that if α is maximal then $\uparrow\alpha$ is consistent and complete.

To prove consistency, assume for reductio that $\uparrow\alpha$ is inconsistent, i.e., that $\bigwedge \uparrow\alpha = \perp$. Since α is a lower bound of $\uparrow\alpha$, we have $\alpha \leq \bigwedge \uparrow\alpha$. So $\alpha = \perp$. But this is impossible because α is maximal.

To prove completeness, assume for reductio that $\uparrow\alpha$ is incomplete, i.e., that $\bigvee \uparrow\alpha \neq \top$. Since α is maximal, $\neg\alpha$ is an upper bound of $\uparrow\alpha$. So $\bigvee \uparrow\alpha \leq \neg\alpha$. It follows that $\neg\alpha = \top$, and hence that $\alpha = \perp$. But again this is impossible because α is maximal. \square

References

- Bergström, L. (1966). *The Alternatives and Consequences of Actions*. Almqvist & Wiksell.
- Bykvist, K. (2002). Alternative actions and the spirit of consequentialism. *Philosophical Studies* 107(1), 45–68.
- Carlson, E. (1999). Consequentialism, alternatives, and actualism. *Philosophical Studies* 96(3), 253–268.
- Castaneda, H.-N. (1968). A problem for utilitarianism. *Analysis* 28(4), 141–142.
- Goldman, H. S. (1978). Doing the best one can. In *Values and Morals: Essays in Honor of William Frankena, Charles Stevenson, and Richard Brandt*, pp. 185–214. Reidel.
- Gustafsson, J. E. (2014). Combinative consequentialism and the problem of act versions. *Philosophical Studies* 167(3), 585–596.
- Jackson, F. and R. Pargetter (1986). Oughts, options, and actualism. *The Philosophical Review* 95(2), 233–55.
- Portmore, D. W. (2015). Maximalism and rational control. Manuscript.
- Shepherd, J. (2015). Deciding as intentional action: Control over decisions. *Australasian Journal of Philosophy* 93(2), 335–351.
- Simons, P. M. (1987). *Parts: A Study in Ontology*. Oxford University Press.