

synopsis of *being for*  
forthcoming in *Analysis Reviews*

My project in *Being For* is both constructive and negative. The main aim of the book is to take the core ideas of metaethical expressivism as far as they can go, and to try to develop a version of expressivism that solves many of the more straightforward open problems which have faced the view without being squarely confronted. In doing so, I develop an expressivist framework that I call *biforcated attitude semantics*, which I claim has the minimal structural features required in order to solve some of these open problems facing expressivism. I take biforcated attitude semantics to prove that expressivism is a coherent and interesting hypothesis about the semantics of natural languages.

So much for the constructive part; having argued that biforcated attitude semantics incorporates the minimal moves required in order to solve a few of the more pressing open questions facing expressivism, I use it in order to productively constrain what an expressivist answer to further open questions must look like. The results, I end up arguing, are not promising; the very same structural features which expressivists need in order to answer so simple a problem as to explain why ‘P’ and ‘ $\sim$ P’ are inconsistent sentences lead to a very general problem about how ordinary, non-moral sentences are to end up with the right truth-conditions, and though I show how to finesse this problem for some simple constructions – truth-conditional connectives and the quantifiers – I ultimately argue that it can’t be done for the full range of constructions in natural languages – including terms like modals, tense, and binary quantifiers like ‘most’. So even if expressivism is coherent and interesting, it is an extremely unpromising hypothesis about the language that we actually speak.

The main theme of the book is that the most fruitful way to understand expressivism is by respecting the fundamental parallel between moral and non-moral language. Expressivism is often introduced by *contrasting* moral with non-moral language, and it is true that at some level, moral and non-moral language *are* different. But the heart of the expressivist view is that the differences between moral and non-moral language are not, at bottom, *linguistic* differences, but rather derive wholly from the differences between moral and non-moral *thought*. Since on an expressivist view, the meaning of every

sentence – moral or not – works by associating some sentence, ‘P’, with what it is to think that P, the only differences between moral and non-moral language are simply derivative from the differences between moral and non-moral thought. Most of *Being For* is really about the consequences of taking this idea seriously – about some of the previously unsolved problems that it solves, and about the new ones that it leaves us with.

The book divides evenly into two halves; in the first half I take on the problems associated with offering a compositional semantics within an expressivist framework for a purely normative language, and in the second half, I use the constraints that are imposed by the solution to these earlier problems in order to attack a further set of problems that arise in an expressivist language with both moral and non-moral predicates. The first half is wholly constructive, and the second half mostly so; it is only in the final chapter that I argue that the constraints that we need to take on in order to solve all of these earlier problems make expressivism an unpromising hypothesis about natural languages like English.

The main problems in the first half of the book are to explain how an expressivist can provide a constructive compositional semantics, and how this semantics can predict and explain some of the simplest semantic properties of complex sentences – such as the fact that negated sentences are logically inconsistent with the sentences that they negate, or that conditionals validate *modus ponens*. A simple argument suffices to show that existing expressivist frameworks – even those with compositional formalisms – do not provide a genuinely compositional semantics. I argue that one simple idea is necessary and sufficient to solve all of the problems considered in the first half of the book. This idea is that rather than each moral predicate semantically contributing a different attitude, we should instead treat all of the sentences of the language as expressing the same very general kind of attitude, and different moral predicates as contributing part of this attitude’s *content*. The name I give to this general type of attitude is *being for*.

So, for example, the sentence ‘murder is wrong’ might express the attitude of being for blaming for murder, and the sentence ‘stealing is better than murder’ might express the attitude of being for preferring stealing to murder. What this view does, is to create a level of structure, the lack of which is what forced previous expressivist theories to postulate infinite hierarchies of distinct attitudes, just in order to allow a single moral predicate to figure in complex sentences. It also allows us to explain all of the properties of classical logic by appeal to one simple assumption about the rational properties of the attitude of being for – that it is what I call being *inconsistency-transmitting*, which just means that two arbitrary people who are in each of two states of being for are ipso facto guaranteed to disagree with one another just in case the contents of those states of being for are inconsistent. (Inconsistency-transmittingness is a familiar property of belief and intention, but not of desire or wondering, for example.)

The only problem with this solution to the problems facing expressivism, is that to work, it requires that all of the sentences in the language express the same underlying kind of attitude. But *prima facie*, the whole point of the expressivist view is to hold that moral and non-moral sentences *don't* express the same kind of attitude. Certainly metaethical expressivists say that non-moral sentences express beliefs, that moral sentences express desire-like attitudes, and that the desire-like attitudes expressed by moral sentences are not a special kind of belief. So the only way for all of the sentences in such a language to all express the same kind of attitude, is for beliefs to turn out to be a special case of the same kind of desire-like attitude – or for both beliefs and desire-like attitudes to be further analyzable in terms of a common, core, attitude that every sentence expresses. The strategy pursued in the second half of *Being For* can be thought of in either of these ways.

What I show in the second half of the book is that we *can* retain the advantages of the first half of the book in an expressivist language with both moral and non-moral predicates, if we have the right underlying analysis of the structure of ordinary, non-moral belief. The semantic framework that takes advantage of this analysis I call *biforcated attitude semantics*, because its underlying idea is that the states expressed by sentences of the language consist of *pairs* of states of being for, one of whose contents is at least as strong as the other's. Biforcated attitude semantics is both conceptually and technically somewhat more complex than the ideas in the first half of the book, but the morals are pretty clear: in biforcated attitude semantics we can again predict and explain all of the right logical properties of complex sentences, and there are several interesting side-consequences, besides. The main problem besetting biforcated attitude semantics has to do with how ordinary, non-moral sentences are to end up having the right truth-conditions, and I offer two alternative solutions to this problem, which I prove to be adequate for languages with expressive power up to that of the predicate calculus.

I think it is important to emphasize that despite the fact that I ultimately argue that it does not work, the constructive framework developed in *Being For* is in no way set up for a fall. Biforcated attitude semantics is both flexible and powerful, and no other expressivist semantic framework ever developed has been both genuinely constructive and formally adequate. In the book I explain why this is, and argue that *any* genuinely constructive, formally adequate expressivist semantics will have the same structural features. A natural response to the book – taken by all three commentators in this symposium – is to look for some place to get off of the boat, before the problems that I discuss in chapter 12 kick in – and the earlier, the better. But my own perspective is different. Despite the fact that the whole point of the book is to draw out the ultimate problems facing expressivism, my own view – which is perhaps a more sympathetic stance than I took at the time I wrote the book – is that these problems remain worth working on, and that the

flexibility and power of the framework provide ample support for the fruitful guidelines that the ideas in the book provide for future theorizing about expressivism.