



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Using Big Words to Explain Little Words

Citation for published version:

Chrisman, M 2011, 'Using Big Words to Explain Little Words', *Think*, vol. 10, no. 29, pp. 23-36.
<https://doi.org/10.1017/S1477175611000054>

Digital Object Identifier (DOI):

[10.1017/S1477175611000054](https://doi.org/10.1017/S1477175611000054)

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Publisher's PDF, also known as Version of record

Published In:

Think

Publisher Rights Statement:

© The Royal Institute of Philosophy, 2011. Chrisman, M. (2011). Using Big Words to Explain Little Words. *Think*, 10(29), 23-36doi: 10.1017/S1477175611000054

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



Think

<http://journals.cambridge.org/THI>


Additional services for **Think**:

Email alerts: [Click here](#)

Subscriptions: [Click here](#)

Commercial reprints: [Click here](#)

Terms of use : [Click here](#)



USING BIG WORDS TO EXPLAIN LITTLE WORDS

Matthew Chrisman

Think / Volume 10 / Issue 29 / Autumn 2011, pp 23 - 36

DOI: 10.1017/S1477175611000054, Published online: 07 September 2011

Link to this article: http://journals.cambridge.org/abstract_S1477175611000054

How to cite this article:

Matthew Chrisman (2011). USING BIG WORDS TO EXPLAIN LITTLE WORDS. *Think*, 10, pp 23-36 doi:10.1017/S1477175611000054

Request Permissions : [Click here](#)

USING BIG WORDS TO EXPLAIN LITTLE WORDS

Matthew Chrisman

Background

Sometimes, when I go to dinner parties organized by my partner, people ask me what I do, and I say that I'm a philosopher. But when I fumble at their questions about 'my philosophy', my partner will describe what I do by saying, 'He uses big words to explain little words.' Although this is meant tongue in cheek, it's basically right. My philosophical research is mainly in metaethics and the philosophy of language with a focus on the semantics of moral words. This means, for better or worse, that I use big words to explain little words.

The specific little words that I study are those used to talk about ethics. Words like 'right', 'wrong', 'good', 'bad', and especially 'ought' – as in 'You ought to tell the truth.' These are words we use all the time and so clearly understand. So, why do we need to study these words? Because there are many different levels of understanding. We all understand these words in the sense that we can use them, but a language teacher who wanted to teach these words to her students would need to understand them in a deeper way. Further still, philosophers and linguists try to understand these words in a deeper way, in order to appreciate their logical connections to other words and concepts, as well as to understand how words like these have meaning in the first place.

The purpose of this paper is to explain why three kinds of theoretical projects require us to use big words to explain little words. Specifically, I will discuss what I see as core projects in 'metaethics', 'philosophy of language', and 'formal semantics'. My plan is to first say something about each of these, then I'll tie them all together in an attempt to

doi:10.1017/S1477175611000054
Think 29, Vol. 10 (Autumn 2011)

© The Royal Institute of Philosophy, 2011

say in a way that's accessible – even if not quite dinner party material – why in my research I use big words to explain little words.

Metaethics

Here are two actions that might be thought to be morally wrong: (i) stealing someone's lunch money for kicks, (ii) condoning a non-defensive military action. These are just examples; you can supply your own if you don't like mine. But whatever the examples, almost everyone has some sense of what actions are right and wrong.

This is the basic subject matter of ethics. We study ethics as a way of refining and grounding this ethical sense, with the hope of being clearer and more justified in our understanding of what's right and wrong. What do ethically wrong actions have in common? When we disagree or are unsure about what we ought to do in a particular situation, how do we settle the issue? The study of ethics is concerned with these kinds of issues.

There may be several ways to study ethics, but I think the *philosophical* study of ethics is fruitfully divided into three branches. *Normative ethicists* seek to articulate general ethical principles (or 'norms'). Here, some of the classical theories are Kantianism and Utilitarianism. From here, philosophers move to something either more concrete or more abstract. At the more concrete end, *applied ethicists* seek to answer questions about specific real-world ethical dilemmas. There are debates about the legitimacy of civil disobedience, war, abortion, pollution, etc. By contrast, at the more abstract end, *metaethicists* seek to answer questions about the legitimacy and role of ethical thought and discourse in our overall conception of the world and our distinctive place in it. What that means is that metaethicists ask questions, such as

- Is ethics universal or culturally relative?
- Are there ethical facts?

- If there are ethical facts, how do we come to know them?
- How it is that our ethical commitments can and should move us to action in various circumstances?

The core project of metaethics is to answer these kinds of questions. Instead of asking for principles to tell when something is right or wrong, as we do in normative ethics, or asking whether some particular practice is right or wrong, as we do in applied ethics, in metaethics we ask about what it means to think or say that something is right or wrong. To even articulate a normative or applied ethical viewpoint, one must use ethical concepts and words. Metaethicists are interested in understanding this general phenomenon.

A key part of this is, obviously, our ethical language. This is why studying the language we use to talk about ethical issues is important for understanding the legitimacy and role of ethical thought and discourse, which is the core project of metaethics.

This is generally how metaethicists (including myself) get interested in the meaning of words like 'right', 'wrong', 'good', 'bad', 'ought'. And almost as soon as one asks about the meaning of these words, one is struck by the central philosophical problem that I have spent most of my (still short!) career working on. On the one hand, ethical words seem to be just like many other words that we use to talk about the world around us, but on the other hand they seem to perform a practical function importantly different from a lot of other words.

To see how they are similar notice that just as I can also say 'The death penalty is wrong' to assert my ethical views about the illegitimacy of the death penalty, I can say 'The death penalty is used in few countries' to assert my non-ethical views about the prevalence of the death penalty. Because of this, it's natural to think each of these assertions expresses one of my beliefs. Moreover, I can, at least

in principle, be called on to give reasons for my ethical and non-ethical assertions; what justifies me in believing that they are true? I can also use both ethical and non-ethical words in similar logically complex constructions. For instance, I can say 'If the death penalty is used in few countries, then it will die out soon' but I can also say 'If the death penalty is wrong, then it will die out soon'. I can formulate related questions 'Is the death penalty used in few countries? Is the death penalty wrong?' In all of these respects ethical language seems to be on a par with non-ethical language.

Because non-ethical language often seems to be describing things, this has encouraged a *descriptivist* theory of the meaning of ethical language. The core idea is that both ethical and non-ethical statements are descriptions of reality. As such, they can be true or false, and they're naturally interpreted as expressing a speaker's beliefs.

However, to appreciate the other side of this debate, notice that ethical language seems to play a distinctive role in our reasoning and talking about action. For instance, if someone says that the death penalty is more costly than the alternatives, we don't necessarily expect her to act in any particular way. Of course, if we know that she is a governmental official primarily motivated by a desire to reducing expense, then we may expect her to oppose the death penalty on these economic grounds. But it's also possible to imagine someone who says this, even though they support the death penalty. Perhaps they think it is right in spite of its cost.

By contrast, if someone says that the death penalty is morally worse than its alternatives, then we *do* expect him to oppose it. This expectation must, of course, be tempered by our understanding of all of the other things that might motivate the individual to do something. Sometimes people do what they say is wrong, because they are insincere, weak-willed, or whatever. If we temper for these, however, the use of moral words seems much more closely tied to motivations for action than the use of other words.

This close link between ethical language and motivation has encouraged some philosophers to reject descriptivism in favor of an *expressivist* theory of the meaning of ethical language. The core idea is that there is an important expressive difference between ethical and non-ethical statements. Proponents of this account say that ethical statements mean what they do because they express some attitude that is different from belief (e.g. a moral sentiment or special sort of preference or desire) which plays a special role in the psychology of motivation. The important thing is that the expression of this attitude is thought to distinguish ethical language from non-ethical language. Indeed, the fact that ethical sentences express this attitude rather than a belief in some fact leads some expressivists to say that ethical statements aren't literally true or false. As the expression of sentiments instead of beliefs about reality, there is nothing objective and independent of us to make them correct or incorrect (although we might of course say that some ethical sentences are true as an especially rhetorical way to express our own moral sentiments).

The similarities and differences between ethical and non-ethical language are important because of the way they lead respectively to descriptivism and expressivism. However, many metaethicists now think that both of these theories are too extreme. We seem to need a middle ground in understanding the nature of ethical language. Otherwise we remain in the dark about the meaning of ethical words, which threatens to infect not only our understanding of ethical thought and discourse but also our understanding of its subject matter – ethics. One way of thinking about much of my own research is that I'm looking for this middle ground.

This is one of the places where I think we need to use big words to explain little words. The little words are the ethical words like 'right', 'wrong', 'good', 'bad', and 'ought'. The big words are whatever theoretical resources we can come up with, in order to clearly articulate and explain the similarities and differences in meaning between these

words and other words. I think the best place to look for these resources is within the philosophy of language. So, I turn to that next.

Philosophy of Language

Let's begin with a contrast between two scenarios: First, the branch of a tree blows in the wind, scratching out an intricate pattern in the dirt. Second, a dying man stranded on an island uses the branch of a tree to scratch out a message that he hopes someone will read after he dies. Although there is a sense in which there may be 'meaning' in both the intricate pattern and the dying man's message, only the latter is meaningful in the sense which philosophers of language try to explain.

They are interested in the peculiar human ability to use complex and arbitrary signs and sounds to communicate with one another. The core phenomenon here is *semantic* meaningfulness: Words have meanings, and when we put them together in the certain ways, we form sentences which have complex meanings. And these complex meanings are related to each other in logically interesting ways. How is this possible?

There are a number of general theories of how words and sentences can have meanings, but at the most abstract level I think we can identify three important ideas. These are not necessarily incompatible, but theories of meaning will differ based on which of these ideas (if any) is thought to be most fundamental.

First, many words and sentences seem to be about things in the world. The words 'chair' and 'table' seem to refer to particular kinds of objects, and the phrase 'is beside' puts into words a relation real-world objects can be in to one another. Thus, the sentence 'The chair is beside the table' can be said to linguistically *represent* something about the way the world is. This kind of connection between words and the features of the world, which they

can be used to represent, has seemed to many philosophers of language to be crucial for understanding the possibility of semantic meaningfulness.

Second, language is clearly an expressive medium (but not the only expressive medium). That is to say that we often communicate by using words to express what's going on inside our minds. When I say 'cat', that expresses a particular idea; and, when all goes well, it evokes in my audience that same idea. Similarly, when I say 'My cat is brown', it's very natural to interpret me as expressing a particular mental state: my *belief* that my cat is brown. When all goes well, if you think I'm being sincere and know what I'm talking about, you'll come to believe that my cat is brown on the basis of my expressing my own belief to the same effect. This is a simplistic example, but the general kind of connection between words and the mental states they can be used to express is a second place many philosophers of language think we need to investigate in order to understand the possibility of semantic meaningfulness.

Third, once words are put together into sentences, they stand in logical relations to one another. For example, the sentence 'My cat is brown' entails the sentence 'My cat is colored'. That is, if it's appropriate to apply the predicate 'is brown' to something, then it's appropriate to apply the predicate 'is colored'. Something about the meaning of 'is brown' and 'is colored' underwrites this logical connection. Moreover, specifically logical words such as 'or', 'all' and 'not' seem to function to encode logical relations between other sorts of words and sentences. The sentence 'All my cats are brown or black' entails that the sentence 'One of my cats is white' is false. Something about the meaning of 'or' and 'all' underwrites this logical entailment. Because of connections like these, many philosophers of language have thought that understanding semantic meaningfulness requires understanding logical relations between words and sentences.

Like I said above, all three of these ideas are important in the philosophy of language. Philosophical theories of

meaning often differ on which of these ideas they take to be most fundamental. The important point here is that the question these theories are attempting to answer is about the foundations of the phenomenon of semantic meaningfulness. We can put the basic question this way: why do words and sentences mean what they do? One place to look for clues for answering this question comes from the attempt by linguists to articulate formally what precisely various sentences of a language like English do mean. This is the subject of formal semantics. I turn to this next.

Formal Semantics

Linguists have long been impressed with two interrelated facts about language. The first fact is that children learn to speak their first language incredibly quickly. If you've ever spent extended time around a toddler who is learning to speak, you will know the phenomenon well. One minute they're learning the word 'twist' for the action needed to open a bottle, an hour later they're applying this word to the action needed to turn on the bath. A week later they seem to know five related words, and a year later they can use them all in sentences with appropriately conjugated verbs. The second fact about language is that most of the sentences we produce and understand in our everyday lives are novel. For most reasonably complex sentences, no one has ever put together these exact words in that exact order to say this exact thing. (That is almost surely true of most of the sentences in this article, for instance.) Yet we understand these sentences perfectly well.

These two facts have lead linguists to endorse as a working hypothesis the idea that language is *semantically compositional*. This means that one should in principle be able to articulate the rules by which the meanings of complex pieces of language (basically, phrases and sentences) are composed out of the meanings of the most basic parts (basically, words) and how they are put

together. No one knows, in anything like a fully comprehensive way, how to articulate the compositional rules of any specific languages, let alone language in general. But semantic compositionality is the working hypothesis of theoretical semantics because it seems to be the only way to explain how young children could possibly learn their native language as fast as they do and adults could continually produce and understand novel sentences with such ease.

The hypothesis is that language users develop an implicit understanding of a relatively limited number of word meanings and compositional rules. Then, they can use these to compose and interpret an incredible number of novel sentences. It must be stressed that this knowledge is thought to be *implicit* know-how. No ordinary speaker is expected to be able to articulate even one of these rules, and fluent speakers clearly don't consciously think about them when they use language. So, it is the task of formal semantics to turn this know-how into explicit principles of a semantic theory. This is much harder than implicitly grasping the rules. (As an analogy, think of what you'd have to write down if you wanted to articulate fully and explicitly how to ride a bicycle to someone who didn't already know anything about balancing, creating momentum, shifting gears, and breaking.)

The basic way formal semanticists try to make the implicit rules of a language explicit is by developing a *semantic model* for that particular language. In order to understand this, it's useful to compare semantics to meteorology.

Meteorologists try to develop a *weather model* for a particular area of the world. This involves a fair bit of abstraction from many of the detailed specifics of the observed weather events. The exact location of the rain might not matter as much as average rainfall in the district. And we might not care about which specific kinds of trees were blown over as much as we care that a quarter of the trees on the north side of all of the streets in a particular town were blown over. Once meteorologists have abstracted

from such details, they can begin to use their model to predict future weather events. The basic aim is to recreate in a different medium and scale something structurally similar to the meteorological laws leading to observed weather events. Once we have a model exhibiting a high degree of predictive success, we can start to use it to understand the nature of the meteorological laws underlying the weather events that we observe.

Similarly, semanticists try to come up with a semantic model for a language. This too involves a fair bit of abstraction from many of the detailed specifics of day-to-day language use. Lots of 'um's and 'ah's are omitted. Idiosyncratic tones and moods are ignored. Idiom and poetry are suppressed. After abstracting from such details, the structures of formal logic are then used to model the content of specific sentences. Linguists do this with the hope of recreating in a different medium and scale something structurally similar to the semantic laws by which the meaning of a whole sentence is a function of the meaning of its parts. The basic aim in semantics is to predict what a sentence will be understood to mean. Once we have a model exhibiting a high degree of predictive success, we can start to use it to understand the nature of the semantic rules implicit in a language, which lead to the creation of novel and complex meanings out of a relatively small number of initial building blocks.

The main goal of formal semantics, then, is to develop a semantic model that is capable of giving us an understanding of the semantic rules of a particular language. This is like meteorology which aims to develop a weather model that is capable of giving us an understanding of the meteorological laws of a particular geographical region. Of course, just as meteorological laws change with time and the evolution of a physical environment, the semantic rules change with time and the evolution of a language. However, we hope this change is slow enough that an accurate snapshot (model) of an evolving system is useful in understanding it.

Putting These Together

I've said something about metaethics, philosophy of language, and formal semantics. I hope this has given you some basic sense of what each of these is about. These areas of scholarship are often pursued independently, but they're each pursued better when one has an understanding of some of the basic ideas in each area.

Initially, I came at this from the point of view of metaethics. I became interested in sorts of metaethical questions bulleted above, which led me to an interest in the nature and legitimacy of ethical thought and discourse. This then led me to wonder about how to explain both the similarities and differences between ethical and non-ethical language. This is the question I presented above: Is the descriptivist or expressivist theory of the meaning of ethical words correct?

What I've concluded is that this question doesn't completely make sense. I now think it runs together two separate but related questions about the meaning of ethical words. The first is a question in the philosophy of language, where we are interested in explaining the nature of the general phenomenon of meaningfulness: Why do ethical words and sentences have the meanings that they do? The second question is a question in formal semantics, where we are interested in developing semantic models to help us understand the compositional rules of particular languages: What is the best way to model the semantic rules of our language, in such a way that it accounts for the similarities and differences in the semantic contribution of ethical and non-ethical words to the sentences in which they figure? The project of answering these questions can be pursued in tandem, but it should be recognized that they're different, and it's not obvious that one or the other is the more fundamental question.

So, why do we need to use big words to explain little words? Let me give two reasons related to what I've

discussed so far and then finish by working through an example meant to illustrate the point.

The first reason we need to use big words to explain little words is that there are a lot of things we can mean by 'meaning'. As we've just seen, there's an important difference – which is often ignored even by the theorists working in this area – between why a word or sentence has the meaning that it does and the rule-governed contribution a word makes to the overall meaning of sentences in which it figures. The former is a question in the philosophy of language, while the latter is a question in formal semantics. The 'big words' like the ones we met above: 'representation', 'expression', 'inferential relations', as well as the formalisms of logic help us to keep these projects distinct.

The second reason we need to use big words to explain little words is that answering these sorts of questions requires conceptual tools that are more precise and fine grained than everyday language provides us. This is not surprising. As we saw above, our ordinary knowledge of the compositional rules of our language is mostly implicit know-how, rather than something that ordinary speakers can articulate explicitly. In order to articulate these rules explicitly, one must develop a vocabulary for talking about a practice (language-use) which we mostly engage in without explicit knowledge of the rules.

Let me try to illustrate these two reasons for using big words to explain little words by considering an example near and dear to me. Consider the English word 'ought'. If we wanted to teach someone who is learning English what this word means and how to use it, we might first try to find a word in their native language that has the same meaning. But not all languages have a word that corresponds directly to the English word 'ought'. For example, in English we can ask about the difference in meaning between 'ought' and 'must', but there aren't separate words for these in French. (Instead French speakers use different moods of the verb to convey this difference.)

So, we might start out by noting that 'ought' is a *modal auxiliary verb*. These are some big words from grammatical theory, but they help in explaining the meaning of 'ought' and how to use it. The fact that 'ought' is a verb means that it plays a specific kind of role in the composition of a sentence. The fact that it is an auxiliary verb means that it must always be combined with some other verb to form the verb-phrase part of a sentence. The fact that it is a modal auxiliary verb in English means, among other things, that it lacks conjugation for person and tense and it forms a system with other words like 'can', 'may', 'might', 'must'. One of the most interesting things about these words is that they have both *deontic* and *epistemic* uses. That means that they can be used to say what is permissible or impermissible according to some set of rules, or they can be used to say what is likely or unlikely according to some set of evidence.

What I've just said about the meaning of 'ought' is just the tip of the iceberg, when it comes to explaining the semantic rules governing the use of the word 'ought'. (I'm currently working on a book which aims to go much deeper.) But this has already required a lot of unordinary concepts and words such as 'modal', 'auxiliary verb', 'conjugation for person and tense', 'deontic', and 'epistemic'. Don't worry if you don't feel like you understand what these big words mean. My point here is that you can have perfectly good implicit knowledge of how to use the word 'ought', but when you start to try to make this knowledge explicit you need a much richer set of conceptual resources even to scrape the tip of the iceberg.

The need to develop a semantic model for ethical words like 'ought' was the first reason I gave above for using big words to explain little words. The second reason has to do with the question of why this word means what it does. Descriptivists would say that this word means what it does in virtue of what parts of the world it can be used to describe. Expressivists would say that this word means what it does in virtue of the non-belief attitude it can be

used to express. I think the fact that this word is a modal auxiliary verb with both deontic and epistemic uses puts pressure on both of these theories. In order to argue that, however, and to develop an alternative requires again the development of new conceptual tools. That is, it requires one to use big words to explain little words.

Conclusion

I have tried here to explain some of the central projects of metaethics, philosophy of language and formal semantics. And I have tried to explain why each of these requires us to use big words to explain little words. We need the big words to keep these theoretical projects distinct, and in each case big words give us new conceptual tools for thinking about how to pursue these projects. I think this justifies the good use of big words to explain little words, but let me close with a warning. As in most areas of scholarship, it's tempting in pursuing these projects to throw jargon at things we don't understand very well. This is the bad use of big words, and it's bad precisely because it doesn't explain anything. In trying to explain things like what it is in virtue of which 'ought' means what it does, or what contribution this word makes to the meaning of the sentences in which it figures, I think we need to be vigilant against the bad use of big words, while recognizing the usefulness of the new conceptual tools represented by the good use of big words to explain little words.¹

Matthew Chrisman is a Lecturer in Philosophy at The University of Edinburgh. matthew.chrisman@ed.ac.uk

Note

¹For helpful feedback on earlier drafts of this paper, I'd like to thank Jean Beckley, Nicole Breazeale, Graham Hubbs, Tara Lloyd, and an audience at the University of Cape Town. Research for this article was supported by the Arts and Humanities Research Council.