

Knowing One's Mind *William Child*

1. Self Knowledge and the Propositional Attitudes

How do we know the contents of our own minds? That question prompts another; what *are* the contents of our minds?

Philosophers tend to divide mental phenomena into two broad categories. On the one hand, there are *experiences*. This category includes at least perceptual experiences (such as the visual experience one has when one sees a tree) and bodily sensations (pains, toothaches etc.). On the other hand, there are such states as *believing* that Oslo is in Norway, *wanting* Senegal to win the World Cup, *intending* to make a chocolate cake. We can think of these states as *attitudes* towards *propositions*. Believing that Senegal will win is taking the attitude of *believing* towards the proposition *that Senegal will win*; wanting Senegal to win is taking the attitude of *wanting it to be the case that* towards the proposition *that Senegal will win*, and so on. So we can call such states *propositional attitudes*. No doubt the division of mental phenomena into just two categories is too crude. But, provided we treat it with caution, it is a helpful way of carving things up.

My topic is our knowledge of our own propositional attitudes. How do we know what we believe, want, intend, and so on? All of us have many beliefs about ourselves. For example, I believe that I was born in Cambridge and that I am over 5'10" tall. And I not only believe these things about myself; I know them. But notice two features of these cases. First, though I do, in fact, know where I was born and how tall I am, we don't think that every rational person must know where they were born or how tall they are; it is easy to think of reasons why someone might not know such facts about themselves. Second, I do not have any *special way* of knowing these facts about myself. Any way of knowing my birthplace or my height that is available to me is also, in principle, available to other people. The way I know where I was born, for instance, is by looking at my birth certificate and asking my parents. And you can know where I was born in exactly the same way.

Now consider our knowledge of our propositional attitudes. People are generally right about what they currently believe, want and intend. Of course, we can be wrong about our true beliefs, desires and intentions. But such mistakes are the exception, not the rule. Being right about one's attitudes is the normal situation.

This sort of self-knowledge differs in important respects from knowledge of such facts about oneself as one's height or birthplace. First, being right about one's own attitudes is the norm. It is not just that people do, as a matter of fact, tend to be right about their beliefs and desires. It is an *essential truth* that people are by and large right about what they believe, desire and intend, in a way in which it is not an essential truth that people are right about their height or birthplace. While it is easy to see how someone could be wrong about how tall they are or where they were born, it is not at all easy to see how someone could, quite generally, be wrong about what they believe,

desire and intend. In fact, you could not be a rational agent at all if you did not by and large know what you believe, what you want, and what you intend to do.

Second, each of us has a special way of knowing about our own propositional attitudes - a way in which no-one else can know about them. If you know that I believe Henman will win Wimbledon, we can always ask *how you know*. And the answer is obvious: your knowledge of what I believe is based on what I say and how I behave. But consider the question, "How do I know that I believe that Henman will win?" In the normal case, my knowledge of what I believe is immediate and effortless; it is not based on anything at all. It seems that I just *do* know what I believe, without even trying. No-one else can know of my beliefs with the same effortlessness and immediacy. The same is true for my desires and intentions.

How are we to understand this knowledge of our own propositional attitudes? I will review three philosophical accounts of self-knowledge, all of them for one reason or another unsatisfactory. Then I will describe a different approach, which looks more promising

2. Three Accounts of Self-Knowledge

There are three popular ways of explaining our knowledge of our own propositional attitudes.

First, there is the idea that all self-knowledge is based on introspected experience - on feelings, sensations and the like. On this view, there is a particular kind of experience associated with each kind of propositional attitude. So there is a particular way it feels like to *believe* that Henman will win Wimbledon, a different way it feels like to *want* him to win, and so on for every different kind of attitude one might have. The source of my knowledge of what I believe and desire, then, is essentially the same as the source of my knowledge that I am in pain or that I have pins and needles. Introspection tells us about our experiences and feelings. And, since belief and desire are correlated with specific kinds of experience, introspecting our experiences and feelings tells us what we believe and desire.

But this account seems hopeless. It is simply not true that there is any distinctive experience involved in believing that such-and-such is true, or wanting it to be the case that so-and-so. Try, for example, to reflect on what it *feels like* to believe that Oslo is the capital of Norway, or that $e=mc^2$, or that one's forename is such-and-such. It is hard to give much credence to the idea that there is something *experiential* that captures what is common to all cases of believing. The same goes for the other propositional attitudes too.

A second suggestion is that ascribing attitudes to ourselves involves a kind of self-interpretation. On this view (contrary to what was claimed in part 1 above) the process of self-ascribing attitudes is no different *in kind* from the process of ascribing attitudes to other people. I observe my behaviour and ascribe to myself the beliefs, intentions and desires that make best sense of that behaviour, just as I ascribe to other people the beliefs and so forth that make best sense of their behaviour. My beliefs about my own attitudes are more likely to be true than my beliefs about other people's (or their beliefs about mine). But the reason for that is just that I have very much more

experience of my own behaviour than I have of other people's behaviour. And that is a purely *quantitative* difference; there is no *qualitative* difference at all between ascribing attitudes to oneself and ascribing them to others.

Now such a process of self-interpretation does have a part to play in self-knowledge. We do sometimes take up a third-person attitude to our own behaviour, *working out* what we believe or desire in something like the way we work out what others believe and desire. But that is obviously not what we do in most ordinary cases of self-ascription. In most cases there is no question of observing our own behaviour. I just know, *without any basis at all*, that I believe that Oslo is in Norway, or that I intend to brush my teeth.

A third suggestion is that our self-knowledge is the product of a reliable causal mechanism. On this view, each person's mind contains a mechanism that, given an intention (for example) as input, produces the belief in the subject that she has that intention. We do not have to *think about* our intentions or behaviour in order to work out what we intend. It is just a feature of the human mind that, when one's mind contains the intention to brush one's teeth (say), that intention automatically causes the belief that one intends to brush one's teeth. So the reason our beliefs about our own attitudes are generally right is that the causal mechanism that produces those beliefs is a very reliable mechanism.

There is something right in the idea that our attitudes reliably cause beliefs that we have those attitudes. When I believe that I intend to brush my teeth, there must be *something* that causally explains my having that belief; it did not come from nowhere. And the fact that I do intend to brush my teeth presumably plays some part in producing my belief that I have that intention. But even if it is *true* that our attitudes cause beliefs about them, it does not follow that the appeal to a reliable causal mechanism is by itself enough to give us a complete account of our knowledge of our own attitudes. In particular, the causal account itself says nothing about what is actually involved, *from my own point of view*, in forming a belief about what I intend, expect or believe. The account seems to suggest that such beliefs about our own attitudes just pop into our minds - that I simply find myself, for no reason, believing that I intend to brush my teeth, or believing that I believe that Oslo is in Norway. And that is untrue to the actual experience of forming beliefs about our own attitudes. In real life, we do not just *find ourselves* possessed of beliefs about our attitudes - as if from nowhere. We have reasons for holding the beliefs about our attitudes that we do; and there are ways in which we reach those beliefs. Without an account of what self-ascription actually involves from the subject's own perspective, the causal account is incomplete.

3. A Different Account of Self-Knowledge

I shall describe a different view of the self-ascription of propositional attitudes, an account that is rooted in suggestions made by Wittgenstein and that has been developed by others in recent years.¹ Consider first the case of self-ascribing beliefs. What is a belief? Seen from your point of view, having a belief is not a matter of being in some internal state. As far as you are concerned, what you believe is simply a matter of how the world is. Suppose you believe that Oslo is in Norway. What that

means is that, from your point of view, Oslo is in Norway. What follows from this apparently uninteresting truism?

Suppose I make the judgement, "Oslo is in Norway". The topic of that judgement is Oslo and Norway; the judgement is not about me. But the fact that I make that judgement does imply something about me. If you know that I have judged that Oslo is in Norway, you know that, *from my point of view*, Oslo is in Norway. And if you know that, then you know that *I believe that* Oslo is in Norway. So you can learn something about me from the judgements I make about the external world. In Wittgenstein's words:

The language-game of reporting can be given such a turn that a report is not meant to inform the hearer about its subject matter but about the person making the report.

It is so when, for instance, a teacher examines a pupil. (PI pp. 190-1.)

But suppose I want to state explicitly that I believe that such-and-such. My judgements about the external world automatically *express* what I believe about the world. But to state that I believe that such-and-such, I must move from a judgement that *expresses* my belief to a judgement that *explicitly* ascribes that belief to me. For example, I must move from the judgement "Oslo is in Norway" (which *expresses* my belief that Oslo is in Norway) to the judgement "I believe that Oslo is in Norway" (which *explicitly states* that I have that belief). How do I do that? Simply by prefixing the judgement I am prepared to make about the external world with the clause, "I believe that . . .". That clause converts my judgement about the external world ("Oslo is in Norway") into a judgement about myself ("I believe that Oslo is in Norway"). And I reach that judgement about myself without the need for any introspection or self-observation. As long as I understand the words "I believe that . . .", there is a simple procedure for ascribing beliefs to myself. To tell what I believe about where Oslo is, for example, this is what I have to do. First, consider the question, "Where is Oslo?" Second, answer that question - by judging (e.g.) "Oslo is in Norway". Third, prefix that judgement with the clause, "I believe that . . .". That is all I need to do to reach a belief about what I believe.

What does this account explain? First, it explains the reliability of our beliefs about our own beliefs. The judgement I make about where Oslo is is already an expression or manifestation of my current belief about where Oslo is. So a modification of that judgement is all that is needed to produce a correct self-ascription of the belief. Crucially, since I form that belief about my belief without examining any evidence about what I believe, there is no room for error to slip in because the evidence is incomplete, or because I make mistakes in assessing it.

Second, the account addresses the question of how, from the point of view of the subject, beliefs about her own beliefs are reached. (So it fills the gap we noticed in the causal account of self-knowledge.) I do not just find myself possessing beliefs about what I believe; I *reach* those beliefs, by considering how things objectively are, understanding that the way things are from my point of view just is the way I believe them to be, and making the simple manoeuvre that turns a judgement that expresses what I believe into a judgement that explicitly ascribes that belief to me.

4. *Some Questions About The Account*

The position I have sketched in section 3 offers a promising account of our self-ascriptions of our current beliefs - an account that describes how we make those self-ascriptions and explains their reliability. But even if this account is successful, it is only the beginning of a full understanding of our knowledge of our propositional attitudes. I will conclude by mentioning some issues that need to be faced.

First, there are instances in which the judgements we make about the world do not express what we *really* believe: cases of self-deception, wishful thinking and the like. Suppose that, in my enthusiasm for an African victory, I judge "Senegal will win the World Cup". Prefixing that judgement with the clause "I believe that . . .", I self-ascribe a belief: "I believe that Senegal will win the World Cup". But suppose that, deep down, I do not really believe that they will. In this case, the model I have described has produced a false self-ascription. So the model, however reliable, is not infallible. Does that show that the model cannot, after all, give me *knowledge* of my own beliefs? And, in cases of self-deception, wishful thinking and the like, how do we know what we *really* believe; what alternative method of self-ascription do we have?

Second, how far can the model I have sketched be generalised? The central idea is that our effortless self-ascription of current beliefs can be explained in terms of the simple conceptual manoeuvre that converts a judgement that expresses a belief into a judgement that self-ascribes that belief. But belief is just one kind of propositional attitude. There are many others. There are attitudes that are quite closely allied to belief: expecting, suspecting, anticipating, etc. And there are others, like desire and intention, that are very different. To give an account of self-knowledge for these cases that parallels the account we have offered for belief, we will need to find, for each case, a kind of judgement that expresses the relevant attitude. And we will have to show that it is possible to convert that judgement, by some simple conceptual manoeuvre, into a judgement that self-ascribes the attitude. But can this be done? *Are* there kinds of judgement that express our desires, intentions etc.? And if so, can they be easily converted into explicit self-ascriptions?

Third, I have spoken only about our knowledge of our *current* propositional attitudes. But what of our ability to remember what we believed or desired or intended in the past? Are we reliable about our own past propositional attitudes? If so, what is the source of that reliability? And can the model that seems to work for present-tense self-ascriptions of belief be extended to the past-tense case? It would not be surprising to find that the model I have sketched cannot be extended to all cases, or to all types of propositional attitude. In philosophy, over-generalising a good idea is a common mistake, and there is no reason to expect that our general reliability about our own propositional attitudes will have a single source. But we should not be too pessimistic about the account described in section 3. The fact that a given account is not correct for all cases does not show that it is not correct for any case.

William Child
University College, Oxford

¹ See Ludwig Wittgenstein, *Philosophical Investigations*, (Oxford: Basil Blackwell, 1953) part II section x. For the classic modern development of Wittgenstein's suggestion, see Gareth Evans *The Varieties of Reference*, (Oxford: Oxford University Press, 1982) ch. 7, especially pp. 225-6.