Knowledge through Communication:

A response to the question of how testimony-based knowledge is possible.

by Thérèse Dewhurst

Thesis Presented for the Degree of

DOCTOR OF PHILOSOPHY

in the Department of Philosophy of the

UNIVERSITY OF CAPE TOWN

December 2014

Abstract: *The aim of this thesis is to offer a response to the question of how it is that a hearer can get knowledge from testimony. The project has two main components. The first is to suggest that the obstacle to getting knowledge through testimony (the obstacle of epistemic vulnerability) is one that can be ignored. The second is to set out how it is that mere communication could be sufficient to explain how testimony can be a source of knowledge.*

The first component constitutes a proposal to reject the problem of testimony as it is usually conceived. Testimony is often seen to be epistemically distinct and interesting because of the apparent epistemic vulnerability posed by its being an indirect source of knowledge. Viewing the problem in this way has led most epistemologists to set out on a project of justification: the challenge is to explain how it is that hearing an assertion can be sufficient grounds for coming to know what is asserted. Whether one is a reductionist or a non-reductionist, the aim has been to establish that essential link between hearing a speaker assert that p, and p's being true. I will argue that seeing the problem of testimony as one of epistemic vulnerability is only inevitable if one has a particular view of knowledge. If we take knowledge to be a state metaphysically distinct from belief, a state not dependent on its justification to establish it as knowledge, then the indirectness of testimony does not inevitably result in a problem of vulnerability.

The second component constitutes a positive explanation for the possibility of knowledge through testimony. I argue that (rather than seeking to justify our testimony-based beliefs) we ought to try to understand the mechanism whereby knowledge can be made available to a hearer simply by understanding an assertion. In this endeavour, I propose a certain theory of communication, such that understanding a communicative utterance entails coming to recognise the speaker's actual mental state. If successful communication gives the hearer access to the speaker's actual mental state, then successful communication can explain how understanding an assertion that p can get a hearer to know that p. I argue that correctly understanding an assertion that p entails coming to know that p*. I defend the idea that the institution of communication explains how knowledge through testimony is possible.*

For Dan, a constant help, and Leo, a constant hindrance

Acknowledgements

First and foremost, I would like to thank my supervisors, Professor Bernhard Weiss, and Dr Dean Chapman, as well as Dr Jeremy Wanderer, for all of their input. Their guidance, criticism and support has played an enormous role in the development of this project. I would also like to thank Richard Flockemann for taking the time to talk to me about my thesis. I would like to thank my parents for all their support and encouragement, and my husband, Dan, for his endless patience and insight.

Contents:

Introduction

We get knowledge from other people telling us things. It is such a common place occurrence that I only have to think of my most recent interaction with another person to come up with an example. At the moment, this is my friend Bo telling me that the bees nesting in the eves have started to swarm. When she told me 'The bees are swarming,' I then knew that the bees were swarming. I couldn't see this for myself, I came to know it based only on hearing her tell me. The aim of this thesis is to explain how this happens: how is it possible to get knowledge from testimony?

This question is one that has received quite a lot of attention in the last twenty or so years. There have been many different ways of formulating the question and many different kinds of response. Explaining how we get knowledge from testimony is usually seen to be a problem because of the apparent incompatibility between the satisfactory epistemic status of the beliefs that we get from testimony (i.e. knowledge) and the prima facie flimsiness or fragility of the basis for those beliefs. I know that Mount Vesuvius erupted in AD79 and covered Pompeii in lava and ash, I know that rhinos are animals threatened by extinction due to perpetual poaching, I know (knew) that the bees are (were) swarming, and that my brother hates his job in Hong Kong. I *know* these things, and yet I know them only because I have been told them. But what kind of basis is that, given that the speaker who told me could be lying or mistaken? How can I know that these things I have been told are *true*?

When a person knows something, that person cannot be wrong. Knowledge is a very strong (the strongest) epistemic position one can be in. How can we have that kind of confidence in beliefs formed on the basis of another person's testimony? What is the *justification* we have for taking another person's word or it? When we get knowledge from perception or inference we do not have to rely on another person's reason or reliability. When I am told by a speaker that something is the case, the possibility that the speaker is wrong, that she is somehow mistaken, or that she is lying, are possibilities that threaten or undermine the epistemic status of my testimony-based beliefs. The indirectness of testimony results in an epistemic vulnerability that must be overcome. This is the problem of testimony as it is usually conceived.

But this is not how I see it. The problem, as I see it, is not about showing how our testimony-based beliefs can have the satisfactory epistemic status that they do, rather it is about showing how human communication can function as a vehicle for knowledge transmission. My account is distinct in that I want to approach the question of how testimony-based knowledge is possible from within a controversial paradigm – a paradigm that rejects the idea that knowledge is a state composed of mental and non-mental components. It is my contention that the problem of epistemic vulnerability is inevitable only if one is committed to the view that knowledge is reducible into something like justified true belief.[1] Rejecting this view of knowledge allows us to avoid the problem of vulnerability: epistemic vulnerability is no longer an inevitable consequence of indirectness. And if epistemic vulnerability is no longer the problem of testimony, the problem becomes one of *explaining how it is that testimony functions as a means to communicating knowledge*.

It is my contention that taking on an externalist conception of knowledge, a conception defended by Timothy Williamson in 'Knowledge and its Limits', can serve to undermine or deracinate 'the problem of testimony' as it is usually conceived. If we do *not* see knowledge as a belief with a special status established on the strength of its justification, then the problem of testimony need no longer be about establishing testimony as being a good or reliable source of belief. Embracing the externalist view of knowledge frees us from having to establish the possibility of testimony-based knowledge by resorting to stories about justification. Knowledge is a mental state identified by its factiveness (its being true) and not by the strength of the reasons or entitlements that we have for believing it. The possibility of getting knowledge through testimony is not threatened by the necessary indirectness of testimony as a source.

Once we are free from seeing the problem of testimony as one of epistemic vulnerability, the question still remains, how is testimony-based knowledge possible? How is it that hearing a person make an assertion, or use words or sounds or gestures, can put

[1] Both Michael Welbourne (2001) and Sanford Goldberg (2010) come to the similar conclusion that understanding the concept of knowledge (and rejecting standard conceptions) is key to understanding how knowledge can be transmitted through testimony. Welbourne defends the idea that knowledge is necessarily communal, which involves a rejection of the standard internalist view of knowledge; Goldberg argues that an 'anti-individualistic' conception of knowledge is required in order to understand how we can get knowledge from testimony, though he maintains that knowledge is analysable into *reliable* true belief.

the hearer in a position to know that the proposition asserted is true? I will argue that it is the institution of human communication that explains how it is that knowledge is transmitted from a speaker to a hearer. When a speaker correctly asserts that p, she must know that p, and when a hearer correctly understands her assertion, he also comes to know that p. A successful instance of testimony *entails* that the hearer come to know what the speaker asserts. There is no epistemic gap between understanding an assertion that p, and coming to know that p is true: understanding it properly *is* coming to know that it is true.

The success of this explanation of how testimony-based knowledge is possible depends (as much as the rejection of seeing the problem of testimony as the problem of vulnerability) on embracing an externalist conception of knowledge.[2] The externalist conception of knowledge is revolutionary. It seems, from my perspective, to be gathering some momentum, but it is still highly controversial and goes against some of the strongest convictions held by philosophers since (and perhaps before) Descartes. The externalist conception of knowledge entails a rejection of the favoured view that knowledge is a state that is analysable into internal (mental) and external (non-mental) components. The idea that truth is external to the mind is almost non-negotiable for some. However, it is exactly this premise that the externalists reject. Instead, knowledge is understood to be a mental state, a *purely* mental state, and yet it is also a state that encompasses the external world.

It is my optimistic opinion that we are witnessing the beginning of a Kuhnian revolution in epistemology. There have been problems plaguing the broadly accepted analysis of knowledge for some time, problems that just won't go away. The broad model is quite generally accepted, but there is very little agreement within this model of what exactly knowledge is: finding an intuitively satisfactory notion of knowledge within this framework is difficult. On the one hand, knowing something entails that what is known is true, and on the other, our faculties of reason and perception are not infallible. There is no easy answer here – how can we know something is true at the same time as allowing that we could be wrong? My hope is that this kind of problem will eventually lead epistemologists to reject

[2] Not to be confused with externalist accounts of *justification*, such as reliabilism. These accounts still fall within the mainstream internalist view that knowledge is reducible into internal and external components. I employ the term internalist in the way that Williamson does: 'Internalism is the claim that all purely mental states are narrow [are determined by internal physical states]; externalism is the denial of internalism.' (p.52) Internalism about knowledge relegates truth or factiveness to the non-mental realm. This has the consequence that reliabilists (externalists about justification) count as internalists.

the idea that knowledge is a state analysable into internal and external components, and will look more favourably on the alternatives.

However, idealistic musings aside, I will use this controversial conception of knowledge to defend what I believe to be a more satisfactory response to the epistemological problem of testimony. The fact that I will be working with a controversial view of knowledge means that much of what I say might seem unpalatable to the reader, especially one that is ill-disposed towards the externalist premise. But I hope that such misgivings can be temporarily laid aside in order that the reader may appreciate the strengths of my account. My aim in this thesis is not to defend the externalist view of knowledge, and I cannot hope to convince a committed internalist to renounce their position. What I do hope to do is to shed some light on one of the possible advantages of making such a renunciation. I propose that if we reject an internalist conception of knowledge (the idea that knowledge is analysable into internal and external components), the possibility of testimony-based knowledge could be satisfactorily explained.

**

The structure of the thesis will be as follows: In Chapter 1 I will set out the problem of testimony and how I intend to respond to it, using a framework suggested by Quassim Cassam in 'The Possibility of Knowledge'[3]. In Chapter 2 I will give a brief overview of the different approaches to testimony so far, and give reasons for thinking that the responses offered are not completely satisfactory: they are united in that they aim to overcome the problem of vulnerability with justification, and yet, I will argue, this is problematic. In Chapter 3 I argue that if we adopt an externalist conception of knowledge, then we do not *need* to take the route of justification in order to show how testimony-based knowledge is possible – the problem of vulnerability is no longer inevitable. The fourth and fifth chapters constitute my proposed substantive explanation for the possibility of testimony-based knowledge. More than just defending the claim that testimony-based knowledge is *not impossible*, I aim to offer an explanation for how it is that knowledge is transmitted through communication. In Chapter 4, I defend an account of successful communication that would entail a particular *factive* kind of understanding on the part of the hearer, and in Chapter 5, I

[3] 2007.

defend the claim that the proper understanding of an assertion (a testimonial communicative utterance) entails recognising that what the speaker asserts is true.

As Michael Welbourne has put it, the fact that knowledge is transmitted through testimony can appear either 'practically banal' or 'quite shocking',[4] depending on one's epistemic commitments. It is my aim in this thesis to move away from the idea that knowledge transmission through testimony is shocking. Rather I want to propose that the fact that we can get knowledge in this way is perfectly banal, given a certain conception of knowledge, and given a certain conception of communication. We are able to share and gain knowledge via testimony because we are able to communicate.

[4] 1979 p.1.

Chapter One: The how-possible question

One thing that I am *not* going to argue for in this thesis is that we get knowledge through testimony. This is something that I will take for granted; that we get knowledge in this way is one of the grounding presumptions of the thesis. I know, for example, how long it takes light to reach the earth from the sun, I know that in 1492 Columbus landed in the Americas, I know that my friend went to Argentina for her honeymoon – things that I learnt from the testimony of others. I take it for granted that I know these things and I will not try to defend this fact against the possibility of scepticism in this thesis. This is not to say a sceptical view (the view that knowledge through testimony is not possible) is not a defensible position, only that the purpose of this thesis is not to argue against it. I will not take the sceptical view as a serious alternative. The purpose of this thesis is to answer the question: how is testimony-based knowledge possible?; not: *is* testimony-based knowledge possible? In other words, my question is: given that testimony-based knowledge is possible, how does it come about?

So, what is the best way to go about answering a how-possible question? Quassim Cassam has written a book called 'The Possibility of Knowledge',[1] and in it he defends, and makes use of, a certain structured type of response to a how-possible question, what he calls a 'multi-levelled response.'[2] In this thesis I will make use of Cassam's suggested approach to epistemological 'how-possible' problems, an approach which he has applied to perceptual knowledge (Chapters 3 and 4 of his book), knowledge of other minds (Chapter 5), and *a priori* knowledge (Chapter 6). In this thesis I will apply this same framework to the problem of testimony-based knowledge. Using Cassam's proposed structure, I will defend a certain way of seeing the problem of the possibility of testimony-based knowledge, and put forward a way in which it can be satisfactorily responded to.

I will not attempt to engage critically with the approach: it is neat and attractive, and fits well with my overall project. Cassam defends what is essentially a two-pronged response to the how-possible question. The first move is to remove any obstacle to the possibility of the knowledge in question (to show how knowledge of a certain kind is *not im*possible) and the second move is to offer a more substantive explanation (to show what

[1] 2007.
[2] p.9.

it is that makes this knowledge possible –establish the conditions that explain the how such knowledge comes about.) The structure of his approach opens up a space in which one can provide a positive or substantive explanation in responding to a how-possible question. How-possible questions arise as a result of a perceived threat or obstacle to a certain kind of knowledge. But removing this obstacle is not the only job of the philosopher: one can also move beyond this and provide or establish what Cassam calls 'enabling conditions' which form the necessary background upon which the knowledge in question is explained.

In this chapter I will outline Cassam's proposed multi-level response to a how possible question, and then (in Section Two) I will suggest how my project, answering the question of how testimony-based knowledge is possible, will fit in to such a structure. I will not provide much detail on the substance of Cassam's arguments and how he establishes the possibility of perceptual knowledge, knowledge of other minds, and a priori knowledge; instead I will focus on the structure of his approach. It is his multi-levelled framework that I intend to make use of.

Section One: The multi-levelled approach

A how-possible question arises as a result of a perceived conflict between something's being possible, and something else appearing to be an obstacle to that possibility. If something is possible, and yet there seems to be some impediment to that thing's being possible, a challenge arises: an explanation becomes necessary – how is this thing possible?[3] This is a typical kind of problem in epistemology; there often appears to be an obstacle to gaining knowledge of something (for example, knowledge of the external world), and yet we do gain knowledge of it nevertheless. The possibility that I am being deceived by a Cartesian evil demon is a prima facie obstacle to me knowing anything at all about the external world, and yet I do know about the external world. How is that possible?

Cassam recommends a three tiered response to a how-possible question in epistemology. Firstly, (i) one must identify the means by which one comes to have the knowledge in question, secondly, (ii) one must identify and then either overcome or dissipate the perceived obstacle to coming to have knowledge in that way, and thirdly (iii)

[3] Cassam 2007 p.1.

one can identify the 'enabling conditions', which are the conditions under which one can come to have knowledge via the means specified.[4] The third tier is not strictly necessary: once the obstacle has been removed, the how-possible question is essentially answered – the obstacle that gave rise to the question no longer poses a threat. However Cassam argues that the third tier can offer a more satisfactory response because it *explains* how the knowledge in question is possible, it does *more* than just show how the knowledge in question is *not im*possible. What follows is a brief description of each level.

i.) Knowledge can be categorised by content, for example, knowledge of the external world, or knowledge of other minds. A how-possible question may arise as to how knowledge of a certain content is possible, and the first step to answering this question, according to Cassam, is to identify the means by which that knowledge is attained. If the question is 'How is knowledge of other minds possible?' the first step is to identify the means by which we can get knowledge of other minds. A response to a how-possible question begins with means identification.

Cassam's project starts with the broad question: how is knowledge of the external world possible?[5] In order to answer this question he sets about identifying the means by which one can gain knowledge of the external world. One of the ways it can be done is through sense perception – what we learn through our senses. (Another is through the word of others, and that, of course, is the focus of this thesis.) We perceive the world with our senses, and can thereby come to know things about it. Once we have identified one of the means by which we can come to have knowledge of the external world, the question becomes: how is knowledge *by this means* possible. For Cassam: how is perceptual knowledge possible?[6]

ii.) How-possible questions are 'obstacle dependent.'[7] What this means is that they arise as a result of, or in response to, a perceived obstacle. The question of how perceptual

[4] *Ibid.* p.9-10. The enabling conditions are necessary conditions for the knowledge in question to come about, but they are not *sufficient*. There will be physiological and biological necessary conditions as well. The enabling conditions that are of interest here are the what Cassam calls 'weakly a priori conditions': the conditions that we, as epistemologists, can establish from the armchair (p.87).
[5] *Ibid.* p.4.
[6] *Ibid.* p.6.
[7] *Ibid.* p.2.

knowledge is possible arises because, although it is admitted that we do get knowledge through our senses, there is a tension between this and the fact that what we get through our senses (sense data/appearances) may appear to be, and yet not be, *true*. Or, as Stroud puts it 'our knowledge of the world is '*underdetermined*' by whatever it is that we get through that source of knowledge known as 'the senses' or 'experience'.'[8] This is a sceptical threat. If it can seem or appear to me that p even when p is not true, then it seems that there must be *more* than just what the senses tell me (how things appear to be) for me to get knowledge about the external world. There needs to be something that gets me from: 'It appears to me that p is the case' to 'p is the case.' Appearing to be a certain way is not enough on its own.

Cassam's suggestion is that there are two ways of responding to this obstacle to perceptual knowledge. On the one hand, the obstacle can be provisionally accepted as genuine: Yes, in order to get knowledge from the senses we need to have something that tells us when what appears to be the case *is* the case. And the project will then be to identify or provide whatever that thing is – and thus to *overcome* the obstacle. So although the obstacle is recognised, it is not recognised as 'insuperable'. But there is another way to respond to the obstacle, and this is to *dissipate* it. Dissipating the obstacle means not taking it on board, denying that it exists.[9] Cassam uses the example of the how-possible question: 'How is evil possible given the existence of God?' to clarify this distinction. The obstacle overcoming approach is to *accept* that God exists, but to argue that evil is possible even if God exists. The dissipating approach would be to *deny* that God exists.

So, in the case of knowledge of the external world being underdetermined by the 'testimony of the senses' (sense data), the two approaches would look like this: To overcome the problem, one would *accept* that what the senses provide us with is not sufficient for knowledge of what appears to be the case, and yet would *deny that this fact is incompatible* with the possibility of getting knowledge through the senses. What would need to be shown is that there is something that gets one from 'It appears to me that p is the case' to 'p is the case'. To dissipate the problem on the other hand, one would *deny* that knowledge of the external world is underdetermined by what we get from the senses. It 'appearing' to me that p is the case *is* sufficient for my knowing that p is the case.

[8] *Ibid*. p.24 (emphasis added).
[9] *Ibid.* p.2.

For Cassam, in his project of responding to the question of how perceptual knowledge is possible, dissipation is the better option. To *overcome* the obstacle of underdetermination, one must be able to show how, despite the fact that what appears to be the case is not sufficient to establish knowledge of what is the case, we can still get knowledge of how things are through our senses. The obstacle of underdetermination makes it the case that what we get through the senses is not knowledge of the world, but knowledge of appearances, and our knowledge of the world is inferred from this 'epistemically prior' knowledge.[10] But, he says, 'inferences from the character of our sensory experiences can't provide us with knowledge of external reality.'[11] If it is accepted that appearances are not sufficient for knowledge, then, according to Cassam, it will be nigh impossible to explain how perceptual knowledge of the world is possible. A better response, it seems to Cassam, is to deny that knowledge of appearances is epistemically prior to knowledge of the world – he denies underdetermination.

Overcoming the obstacle would mean accepting that what we get from the senses is not sufficient for knowledge of the external world, but denying that this is incompatible with knowledge of the external world. In other words, one denies that the obstacle is a genuine obstacle. Dissipating the obstacle means denying the existence of the obstacle, rather than denying its incompatibility with the possibility of knowledge of the external world. Cassam defends the dissipatory response by arguing that knowledge of appearances is not epistemically prior to knowledge of the world; that knowledge of the world through the senses is not inferential. In other words, he denies that knowledge of appearances is insufficient for knowledge of the external world.

I will not go into his arguments here, but what he makes clear (for my purposes) is that an epistemic threat or obstacle to the possibility of knowledge need not be overcome to be removed. If an obstacle to knowledge cannot be overcome, then there is good cause to argue that the obstacle does not exist (unless one is prepared to accept that knowledge of this kind is impossible), and one ought to look for ways to dissipate it instead. Cassam's response to the how-possible question of perceptual knowledge shows that we do not need to meet the sceptic's proposal on his own ground; we do not need to accept his challenge. The sceptic accepts that knowledge through perception is underdetermined by the senses,

[10] Cassam p.26.
[11] *Ibid.* p.26.

and claims that this is an obstacle to knowledge, but this obstacle can be denied if we reject the premise that knowledge through perception is underdetermined by the senses.

iii.) Once the obstacle has been removed (either overcome or dissipated), it is possible for the epistemologist to claim that her work is done. She has established how it is that knowledge via a certain means is not impossible (how it *is* possible in spite of a perceived obstacle to that possibility), and therefore nothing more is required of the epistemologist. However, Cassam argues for a form of what he calls *anti-minimalism*.[12] He argues that the story need not end there, that a satisfactory response to a how-possible question could also seek to *explain* just how it is that knowledge of this type is possible. He calls this moderate anti-minimalism. *Extreme* anti-minimalism would be the idea that a response to a how possible question *must* offer more than just a response to the obstacle. Moderate anti-minimalism is the view that there are ways in which an epistemologist *can* do more to explain how knowledge of a certain kind is possible. And this would be a more satisfying response than a merely minimalist response which entails only the removal of the perceived obstacle.

So level three of Cassam's multi-levelled response proposal is the level at which the epistemologist seeks to establish or explain just how it is that knowledge via a certain means is possible. To do this, she must identify what the *enabling conditions* are. The enabling conditions are the conditions which are in place that explain or make sense of how knowledge of a certain kind is gained.[13] This third level of establishing enabling conditions offers a more substantial or *positive* response to the how-possible question. The enabling conditions are not obstacle removing conditions, they do not play the role of explaining how knowledge can be had in the face of an epistemic threat. They are, more basically, conditions that explain how it is that a certain means to gaining knowledge can be such a means.

The enabling conditions that Cassam seeks to establish are what he calls 'weakly a priori', or conditions which do not require empirical investigation to establish.[14] There will be physiological and biological conditions that are necessary in order to *see that* something is the case, but these conditions are empirically established. Cassam is interested in

[12] *Ibid.* p.36.
[13] *Ibid.* p.9.
[14] *Ibid.* p.87.

establishing a priori, or from the armchair, the conditions that enable one to know something from perceiving it. These conditions are explanatory in that their being in place explains how it is that knowledge can be gained in a certain way. They are not meant to be jointly sufficient: these conditions being in place does not mean that one will always get knowledge, but they explain how it is that one can get knowledge in a certain way. Not all necessary pre-existing conditions are explanatory. Cassam is interested in explaining the possibility of a certain kind of knowledge, without listing all the necessary pre-existing conditions. Enabling conditions are 'a sub-class of necessary conditions', they are 'more specific' than just necessary conditions, they are 'necessary conditions for achieving something by a particular means.' [15]

Cassam embarks on the project of establishing the enabling conditions for perceptual knowledge. In order for one to get perceptual knowledge that, for example, the cup one is holding is chipped, one must be able to *see that* it is chipped. Cassam argues that in order to 'see that the cup is chipped' (an instance of getting knowledge from perception) one must have the capacity for spatial perception, and must also be able to 'understand' what one is seeing, one must have the concept of 'cup' and 'chipped.': he argues that in order to get perceptual knowledge of the external world, one must fulfil the SPR (spatial perception requirement) and the CTR (categorical thinking requirement.) These requirements, when fulfilled, make it possible for one to get knowledge about the world through perception.[16]

Cassam's multi-levelled response to how-possible questions satisfies the minimalist, who is only worried about removing the obstacle to the possibility of knowledge, but it also does more than this. The enabling conditions are conditions that, when in place, can *explain* how perceptual knowledge is possible. They form the necessary background upon which we can get knowledge from our senses. This multi-levelled response to how-possible questions is an approach that opens up a space to explain how knowledge of a certain kind is possible, in a richer sense. It is this that attracted me to Cassam's approach. It seems to me that the most common approach to the problem of testimony is explanatorily minimalist. It is too often treated as a problem of obstacle removal, and not enough attention is given over to *explaining* how knowledge is transferred in this very special way. In the next section I will outline how I intend to fit the problem of testimony into Cassam's multi-levelled response.

[15] *Ibid*. p.17.
[16] *Ibid*. p.88.

Section Two: Testimony and the multi-levelled approach

Cassam began his enquiry asking the question: How is knowledge of the external world possible? He identified sense perception as one of the means by which one can get knowledge of the external world. Testimony, the word of others, is another means to gaining knowledge about the external world. It is also a means to gaining knowledge about the minds of others. I can find out how Marie Antoinette died by being told so, and I can also learn that my sister is feeling stressed because she told me so. I have access to facts about the world, and about other minds, through testimony. The question that I intend to answer in this thesis is: How is testimony-based knowledge possible?

i.) The first level in Cassam's approach is means identification. This project is already a question about means rather than content, but there is a question about what exactly testimony-based knowledge is, and how testimony is distinct from other sources of knowledge. 'Testimony' is often used in a formal sense, as in 'eye-witness testimony', but it is not only this narrow sense of testimony that interests epistemologists. So what exactly is testimony? Most straightforwardly, one gaining knowledge that p from hearing it asserted that p would count as an instance of knowledge through testimony. There is a speaker who makes an assertion and a hearer who understands it, and, all going well, the hearer acquires knowledge of the proposition asserted. But hearing an assertion that p is not always an instance of testimony. This is because assertions are not always intended to be communicative.[17] I can assert 'What a lovely day!' or 'You're late.' without *telling* the hearer that it is a lovely day, or that he is late. And nor are assertions necessary for testimony. I may come out of a job interview and give a thumbs up sign thereby telling my friend that it went well, without, technically, asserting that it went well.

However, the paradigmatic instance of testimony involves a speaker, a hearer, and an assertion. The speaker asserts that p, the hearer understands the assertion that p, and can thereby come to know that p. The concept of testimony becomes much clearer in Chapters 4 and 5, but what distinguishes it as epistemologically interesting and distinct is that it is the

[17] There is a good discussion on the difference between the norms of asserting and the norms of telling by Charles Pelling, 2013.

means by which one gains knowledge of something (that he did not know previously), and gets to know it based on somebody else's intentional sharing of her knowledge. Unlike perception and inference (or reasoning), knowledge based on testimony comes from *another person* and not directly from the fact known. One can independently discover that p from perception or from inference, but one cannot independently discover that p from being told that p. Testimony-based knowledge is *dependent* on the knowledge of another.

Testimony is a special source of knowledge because it involves a transfer of knowledge from one person to another. A person comes to know that p, and then can bring about knowledge that p in a hearer, by telling him that p. The hearer gets knowledge of p not by experiencing or perceiving that p, or working it out that p, but by believing the speaker when she tells him that p. The content of proposition p is first known by the speaker, and then, via testimony (the asserting of proposition p), can be known by the hearer. This makes testimony epistemologically distinct because it is essentially 'indirect'. Other sources of knowledge might be considered more 'direct' because the knowledge is not mediated through the mind of another person. When I see that p, or work it out that p, I am getting knowledge based on my perceiving that p or my reasoning to the conclusion that p. When I get knowledge that p through testimony, I am depending on another person's seeing that p or working it out that p. Sometimes knowledge through testimony is called 'second-hand' knowledge for this reason.

ii.) The second step to answering the how-possible question is the obstacle removal stage. This stage is about overcoming or dissipating the obstacle that generates the question of how testimony-based knowledge is possible. The obstacle to the possibility of knowledge through testimony is the obstacle that its second-hand nature generates. The problem is expressed in a number of different ways, but Jonathan Adler summarises it as: 'An epistemological problem enters … if our ground for coming to these beliefs [testimony-based beliefs] is only the speaker's word, since that seems a very weak basis. What reason, if any, is there for a hearer to just take the speaker's word, given that the speaker is capable of lies, deception, error, and poor, ambiguous, or misleading expression?'[18] The obstacle facing testimony-based knowledge is one of *epistemic vulnerability*. It cannot be the case that the only grounds one has for coming to know what is asserted is hearing and

[18] in Stanford Encyclopaedia of Philosophy (2014).

understanding the assertion: knowledge is something that holds a special, durable epistemic status, and the word of others seems to be such a flimsy basis for it. The obstacle to testimony-based knowledge is that its second-hand nature undermines the epistemic status of the hearer's resultant belief.

The obstacle is not that testimony is indirect, or that a speaker could be lying or mistaken, the obstacle is rather that these facts undermine or challenge the epistemic status of our testimony-based beliefs, challenge their claim to knowledge. The problem is that knowledge or justified belief seems to be threatened by the fact that it is second-hand. Simply understanding an assertion is not, on its own, sufficient grounds for coming to know what is asserted. There needs to be something that links a speaker's assertion that p to the fact that p, something that makes it the case that accepting the speaker's word (believing her assertion) can give us knowledge of what she asserts. There is a parallel with Stroud's claim that knowledge of the external world is *underdetermined* by what we get through the senses. What we get from the senses (sense data) is *insufficient grounds* for knowledge of what is experienced. The vulnerability problem facing testimony is that the knowledge that we get from testimony is underdetermined by simply understanding the assertion.

The overcoming response to this obstacle is to *accept* that getting knowledge of what is asserted requires more than just hearing and understanding an assertion, and to provide whatever it is that is missing. The overcoming response accepts the obstacle, but attempts to overcome it, by arguing that it is not incompatible with the claim that we can get knowledge from testimony. It is to argue that, although the second-hand nature of testimony seems to threaten or undermine the epistemic status of these beliefs, the threat can be overcome by showing how it is that testimony is in fact a reliable source of belief, or that we are entitled to or justified in believing the word of others. The overcoming response is to provide the grounds or the missing epistemic link between hearing and understanding an assertion that p, and coming to know that p.

The dissipatory response, on the other hand, would be to deny that the obstacle exists: to deny that knowledge that p is underdetermined by understanding as assertion that p. The *overcoming* approach accepts that the second-hand nature of testimony threatens testimony-based beliefs, but denies that this is incompatible with getting knowledge or justified belief from testimony. Testimony-based beliefs are justified or candidates for knowledge nevertheless. The *dissipatory* response, on the other hand, denies

the existence of the original threat (that the second-hand nature of testimony implies epistemic vulnerability). The dissipatory response to the problem is to deny that merely understanding an assertion that p is insufficient grounds for coming to know that p.

If the obstacle takes the form: 'Merely understanding an assertion is insufficient grounds (inadequate justification) for coming to know what is asserted,' then overcoming the response means accepting this, and denying it's incompatibility with knowledge by providing or establishing what it is that makes hearing testimony sufficient grounds for knowing what is asserted. Overcoming the obstacle means denying that the fact that understanding an assertion is insufficient grounds for knowledge *results in* the impossibility of getting testimony-based knowledge. It entails providing justification or grounds for accepting testimony, it entails providing the missing epistemic link between understanding the assertion that p, and coming to know that p. Overcoming the obstacle means denying that the supposed threat is incompatible with the possibility of testimony-based knowledge.

The challenge, then, for those wishing to overcome the threat is to come up with a justification/ warrant/entitlement that will give testimony-based beliefs the epistemic status necessary for knowledge. In other words, in order to establish how testimony-based knowledge is possible, the question that must be answered is 'what provides testimony-based beliefs with sufficient grounds/justification to give them the status of knowledge in the face of the threat posed by its second-hand nature?' It is my belief that most approaches to the problem of testimony have been about overcoming the obstacle. The job of the epistemologist here is to establish the justificatory link between understanding an assertion and its being true in order to establish testimony-based beliefs as candidates for knowledge. Overcoming the vulnerability problem is about showing how the fact that understanding an assertion that p is not sufficient to give a hearer knowledge that p is not incompatible with the claim that we can get knowledge from testimony, nevertheless.

The claim that most approaches to the problem of testimony seek to *overcome* the obstacle is a substantive one, and one that I will aim to defend in the next chapter (Chapter 2). I will argue that the three broad areas of response to the problem of testimony (reductionism, non-reductionism and the trust/assurance views) can all be seen to be working towards the same outcome: that of establishing that our testimony based beliefs are justified, or that we are epistemically entitled to them. It is generally accepted in the field of testimony that the job of the epistemologist is to provide testimony-based beliefs

with justification or entitlement that will give them the status of knowledge: the possibility of knowledge is established by showing that, although understanding an assertion is not sufficient to get knowledge of what is asserted, there are other factors (justification/entitlement) in play that can explain how we do get knowledge.

Once I have defended the idea that the different approaches to testimony all aim towards this goal of justification (broadly construed), I will give reasons for thinking that, even if they achieve their aim of justification or entitlement, they will not completely succeed in overcoming the obstacle to the possibility of knowledge through testimony. I will suggest that establishing that testimony-based beliefs are justified, or that we are entitled to have them, is *not the same* as establishing that they count as knowledge. The overcoming approach goes some of the way to allaying worries about vulnerability, but, I contend, does not do enough to show how *knowledge* is possible in the face of the possibility of error. If the obstacle to knowledge is that understanding an assertion is not sufficient grounds for knowing that p, establishing that we are justified in accepting testimony will not (I will argue) be sufficient to overcome this obstacle.

If it is the case that we have reason to think that the obstacle to knowledge posed by its second-hand nature cannot be satisfactorily overcome, then we should (I contend) attempt to dissipate the obstacle instead. Dissipating the obstacle means *denying the existence of the obstacle*, rather than denying *its incompatibility with the possibility of knowledge*. It involves denying that 'Merely understanding an assertion is insufficient grounds (inadequate justification) for coming to know what is asserted,' exists as an obstacle to knowledge. Dissipation will not involve arguing how understanding an assertion can provide sufficient grounds for knowledge, but rather that providing 'sufficient grounds' is not a necessary condition for the establishing of knowledge. My claim will be that the vulnerability problem, the prima facie demand for justification, is one that need not be recognised as a problem. It is my belief that the vulnerability problem arises not just out of the necessary second-hand nature of testimony-based knowledge, but from this fact *combined* with a commitment to an internalist (and distorting) view of knowledge. Rejecting this view of knowledge, I will argue, could allow one to deny that the second-hand nature of testimony results in an epistemic threat to the status of testimony-based beliefs. It is only if one is an internalist about knowledge that one *inevitably* sees an epistemic gap between hearing and understanding an assertion that p, and coming to know that p.

Dissipating the obstacle will be the aim of Chapter 3. I will argue that it is only if we accept the problematic view that knowledge is dependent for its status as knowledge on its justification that we will conclude that its essential second-hand nature results in a demand for justification. In other words, the obstacle to the possibility of testimony-based knowledge (the idea that hearing and understanding an assertion is not sufficient grounds for coming to know what is asserted) arises out of a commitment to a specific conception of knowledge. If we reject the idea that knowledge is reducible to internal and external components, then we will not need to be committed to the idea that the second-hand nature of testimony-based beliefs results in its being epistemically vulnerable. What I advocate is that we should embrace an externalist account of knowledge which releases justification from the burden of establishing beliefs as knowledge.

The inevitability of the epistemic gap, the idea that understanding an assertion is insufficient grounds for knowing that p, arises, it seems to me, as a result of a commitment to the idea that merely believing that p based on hearing it asserted that p is not enough for knowledge because knowledge requires a *proper justification*. The possibility that a speaker could be lying or mistaken means that an assertion that p does not necessarily mean that p is true, and thus, in order to know that p, it must be shown that testimony is a reliable or good source of belief. A belief, in order to count as knowledge, must be justified. Rejecting the conception of knowledge that holds it is a special kind of belief releases one from having to establish justification in order to overcome the epistemic gap. Instead, one could argue that there is no epistemic gap to be crossed. I will ultimately argue that properly understanding an assertion that p *entails* coming to know that p.

iii.) Once I have argued that the perceived obstacle to testimony-based knowledge can be dissipated by showing that the demand it makes on epistemologists need not be met, I will attempt to answer the substantive question of what makes testimony-based knowledge possible. I will attempt to explain how it is that we can get knowledge from testimony by identifying the conditions which make it possible. This is level three of Cassam's structured response: the level of establishing enabling conditions. It is this level of explanation that is not usually offered by epistemologists: once the obstacle has been removed, the reason for the how-possible question arising has been dealt with. However, it is not necessary to stop

there. As Cassam argues, 'there is more to explaining how something is possible than showing that it isn't impossible.'[19]

A satisfactory response to the problem of testimony should *explain how* knowledge through testimony happens. When a hearer gets knowledge from hearing an assertion, in the most straightforward case, he gets knowledge of the content of the proposition that is asserted. A speaker asserts 'p', and the hearer hears her, understands her assertion, and can thereby come to know that p is the case. Phrased like this, knowledge from testimony can almost seem like a magical occurrence. A person emits sounds from her mouth, sounds with a certain meaning, and that meaning is grasped by the hearer, and not only grasped, but believed, and not only believed, but known to be true. And knowing something is to be in a certain relationship with the world, such that what is believed cannot be false. How can all this happen, just by hearing the assertion that p?

The aim at this level is to establish the 'weakly a priori' conditions that must be in place for testimony to be a source of knowledge. The aim is not to justify testimony-based beliefs, it is not to explain how testimony-based beliefs get the privileged epistemic status of knowledge. The aim is not to remove a threat to the possibility of testimony-based knowledge. The aim is to show what enabling condition(s) need to be in place in order for it to be possible that understanding an assertion can get a hearer to know what is asserted. These conditions are not meant to be jointly sufficient – enabling conditions are a 'sub-set' of all the necessary conditions. Their being in place explains how it is possible for hearing an assertion to *be a means* to coming to know what is asserted.

I contend that, in order for this to be possible, it must be the case that when a speaker asserts or 'tells' a hearer that p, the hearer properly understanding her utterance entails his coming to know the proposition that she asserts. If her assertion is correct, and he understands it correctly, then the hearer comes to know what she asserts. The speaker is not just making sounds, she is making sounds with a certain meaning, and she is not just making sounds with a certain meaning, but *doing something* with those meaningful words, and what she is doing with those words, all going well, is *communicating*. Communication is a joint action performed by a speaker (or speakers) and a hearer (or hearers). A speaker makes an utterance, and the hearer understands the utterance. When communication

[19] p.9.

succeeds, the speaker succeeds in getting the hearer to recognise her actual mental state. In Chapter 4 I will defend a conception of communication such that it is an institution dependent upon *constitutive rules* for its existence – one of which is a rule governing how we ought to *understand* utterances. I will argue that when these rules are conformed to, when communication *succeeds*, the hearer is able to recognise the speaker as being in the mental state she is trying to get him to recognise her as being in *and* that this is her actual mental state: properly understanding the utterance requires recognising the speaker's *actual* mental state.

The kind of factive understanding required for communication to succeed is also what explains how knowledge is transmitted through testimony. Assertions, or testimonial utterances, are utterances that, when correct, are attempts to get the hearer to recognise the speaker's knowledge of what she asserts. Properly understanding assertions (testimonial utterances) entails recognising what is asserted *is true*. When a speaker knows what she asserts, what she is telling the hearer is that p is true, and when the hearer properly understands this, what he is recognising is *that p is true*. The hearer having the capacity to communicate is what explains the transmission of knowledge from a speaker to a hearer in an instance of testimony. There is no epistemic gap to cross between understanding an assertion and coming to know what is asserted, because, when all goes well, understanding an assertion *is* coming to know what is asserted. Defending this claim will be the task I undertake in Chapter 5.

The enabling condition that allows for hearing an assertion to be a means to coming to know what is asserted is the institution of communication. This institution on its own is not sufficient to explain how knowledge can be got in this way: other conditions, such as the physical ability to hear, the capacity for conceptual thought, and context dependent conditions will need to be in place for an instance of communication to bring about knowledge in the hearer. The existence of institution of communication does not entail that the hearer can hear the speaker, that he is listening to her, or that he can understand the language she is using. However, when communication *does* occur, these other conditions *must* in place. When the hearer has properly understood an assertion, he was able to hear the speaker, he was listening to her, he could understand the language she used etc. Other conditions may need to be in place in order for communication to occur, but when it does

occur, and when the constitutive rules are being conformed to, the hearer gets knowledge of what is asserted.[20]

Conclusion

The goal of this thesis is to propose an answer to the question: 'How is testimony-based knowledge possible?' There are two dimensions to my response. The first is to argue that there are reasons to take a different kind of approach to the problem to the one that is usually taken, to argue that an approach that aims to *justify* our testimony-based beliefs could be seen to be both insufficient, and unnecessary. It is insufficient, I will argue, because establishing justified belief is not tantamount to establishing the possibility of knowledge. I will argue that it is not necessary because we do not need to see the problem of testimony as one of vulnerability: we can embrace a conception of knowledge that does not lead us inevitably to the idea that testimony-based beliefs require justification in order to establish them as candidates for knowledge. This first dimension is constituted by the next two chapters (2 and 3). The second dimension is to *provide* the different kind of response, a response that explains how testimony-based knowledge is possible without aiming to justify testimony-based beliefs. I will propose an account of communication that shows how it is that a successful instance of testimonial communication *entails* the hearer coming to know what is asserted.

This project fits well with Cassam's structured approach to how-possible questions. He argues that a how-possible question is obstacle dependent, and that there are two possible ways to deal with the obstacle. The obstacle to the possibility testimony-based knowledge is the fact that understanding an utterance is considered to be insufficient grounds for coming to know what is asserted. This problem can be overcome by providing the grounds (justification) that establish testimony-based beliefs as knowledge, or it can be dissipated, by rejecting the idea that understanding an utterance is not sufficient grounds for coming to know what is asserted. My response, unlike most others in the field, is to

[20] As with Cassam's project on perceptual knowledge, listing all the necessary pre-existing conditions is not *explanatory*, per se. The institution of communication explains the possibility of getting knowledge through testimony in a way that the ability to hear, and the ability to speak a certain language, and certain context dependent conditions (though necessary) do not.

dissipate the problem. But Cassam argues that obstacle removal need not be the end of the story: one can establish enabling conditions, which is to say, one can attempt to give a substantive explanation for how testimony-based knowledge can be had. The second dimension of my response fits neatly under this heading: my account of communication is the enabling condition for the possibility of testimony-based knowledge.

The satisfactoriness of my response to the question of how testimony-based knowledge is possible depends upon a rejection of internalism about knowledge, and also a rejection of more popular conceptions of communication. It is not within the scope of the thesis to defend these views against all possible objections,[21] and for this reason my aim is not to *convince* the reader with my arguments, but, more modestly, to make a satisfactory account plausible. I believe that *if* we understand knowledge as an external mental state, and *if* we understand communication in the way that I propose, then we can understand how it is that knowledge can be had through testimony. The question of how testimony-based knowledge is possible can be given a substantive response. I will defend the idea that knowledge through testimony is not threatened or undermined by its second-hand nature (Level 2, the obstacle removal stage), and that it is possible for a speaker's knowledge to be communicated to a hearer via an assertion (Level 3, the positive explanation stage.)

[21] Though, luckily, as regards the rejection of internalism about knowledge, this has already been done for me by Timothy Williamson and John McDowell, amongst others.

Chapter Two: Overcoming the obstacle

In the previous chapter I outlined Quassim Cassam's obstacle dependent, multi-levelled response to how-possible questions, and suggested how I intended my project to follow a similar structure. The question that I intend to answer is: how is testimony-based knowledge possible? The first step is means identification, which I covered in the previous chapter by pinpointing what it is that makes testimony a distinct source of knowledge. What differentiates testimony from other sources of knowledge is that it is necessarily indirect, testimony-based knowledge is second-hand knowledge. A hearer hears a speaker assert that p, and based on this, can come to know that p. Knowledge gained from testimony comes from or depends upon another person.

The distinctiveness of testimony as a second-hand or indirect source has led many epistemologists to question how it is that such knowledge is possible. The second-hand nature of testimony is seen to pose a *vulnerability problem*. Unlike perceiving or inferring that p, merely understanding a speaker's assertion that p seems too weak or flimsy to count as the basis for knowledge that p. Speakers can lie or be mistaken, and thus what she asserts could be false: merely understanding a speaker's assertion is not sufficient grounds for coming to know what she asserts. The problem of testimony, the obstacle that gives rise to the how-possible question, is an *epistemic gap* brought about by its second-hand or indirect nature. The necessary indirectness of testimony as a source of knowledge seems to pose a threat to the epistemic status of our beliefs gained thereby. This is the obstacle to the possibility of testimony-based knowledge.

The second step in Cassam's picture is the removal of the obstacle. One can either accept the obstacle, and attempt to overcome it, by showing that the obstacle is not incompatible with the possibility of testimony-based knowledge; or one can dissipate the obstacle by denying that it exists. If one accepts the obstacle, one accepts that merely understanding an assertion is insufficient grounds for coming to know what is asserted. What this means is that, if one wants to establish that testimony-based knowledge is nevertheless possible, some story must be told to show us how testimony can count as sufficient grounds for knowledge. The obstacle of vulnerability can be overcome if we explain how our testimony-based beliefs are not as vulnerable as they might, at first, appear: the vulnerability problem can be overcome by showing how, despite its second-

hand nature, testimony is a good or reliable source of belief, or that we are warranted in accepting or entitled to accept assertions as true.

The aim of this chapter is, firstly, to defend the proposal that most approaches to solving the problem of testimony so far have been about overcoming the perceived obstacle: they have been attempts to provide justification (or warrant/entitlement) for our testimony-based beliefs, and thereby show that we can have knowledge based on a speaker's word. In other words, their project has been to establish that knowledge through testimony is possible because understanding an assertion is supplemented by some kind of justification or epistemic entitlement that allows testimony-based beliefs to count as knowledge. If there is an epistemic gap between hearing that p from a speaker and coming to know that p, this gap can be bridged if the right kind of justificatory story is provided. This, I will argue, is the project embarked upon by most epistemologists in the field.

The second claim I will try to defend in this chapter is that there is reason to think that viewing the problem in this light, and attempting to solve it via provision of justification, will prove ultimately unsatisfactory if one is attempting to establish the possibility of testimony-based *knowledge*. If the aim is to establish that we can get justified belief from testimony, then this can be achieved; but justified belief is not knowledge. Within the epistemology of testimony, this distinction is often blurred or ignored, but this may be problematic. It is my view that trying to establish the possibility of knowledge by providing justification puts too much pressure on justification: in order for justification to establish a belief *as knowledge* it must establish that the belief in question is true, and this is something that a justification cannot do. What this suggests is that there are grounds to attempt to dissipate rather than overcome the obstacle of vulnerability.

It should be emphasised that the role of this chapter is *not* to argue against the other approaches to testimony, or even to criticise them on their own terms. (The brevity and superficiality of my treatment of these other approaches does not give them a fair hearing.) The aim here is simply to *motivate* for (not necessitate) a different kind of approach to the problem of testimony, a non-justificatory approach. If I can provide any reason to think that a dissipatory response is warranted, then I have succeeded. It need not be the case that the overcoming responses discussed below *are* ultimately unsatisfactory, only that there is a

reason to think that they may be. And if this is so then an approach to the problem which does not involve an attempt to overcome the problem becomes an acceptable alternative.

Section One: The overcoming approach

The perceived obstacle to testimony-based knowledge is one of epistemic vulnerability. This vulnerability arises out of the fact that testimony as a source of knowledge, unlike other sources, depends on another person. In order for the hearer to get knowledge of something from hearing a speaker tell him something, what she says must be true. This means she has to be sincere (not lying) and she has to be competent (not mistaken in what she claims.)[1] The epistemologist's project, if she accepts this obstacle, is to show how testimony is, in fact, reliable, or how a hearer's testimony-based belief is *justified*, or how it is that he is epistemically *entitled* to rely on a speaker's word, or how he can be *warranted* in accepting that what a speaker says is true. These projects attempt to establish the connection between what a speaker asserts, and that assertion's being true. If we get knowledge from testimony (which few deny[2]) then there must be some justification or justificatory story that establishes testimony-based beliefs as knowledge, or as justified beliefs. In order for testimony-based knowledge to be possible, one must establish that we are epistemically entitled to accept what a speaker says as true.

This is the project of the epistemologist who seeks to *overcome* the obstacle of vulnerability. The fact the possibility of knowledge through testimony depends on a speaker being sincere and competent means that there must be something that overcomes or combats this vulnerability. Merely understanding the speaker's assertion is not sufficient, on its own, to explain how we can get knowledge of what is asserted. The way to overcome this obstacle is to establish that, despite the epistemic vulnerability posed by the indirectness of testimony, the justification or entitlement that one gets in an instance of testimony explains

[1] 'Sincerity' and 'competence' are Elizabeth Fricker's terminology; she argues that a speaker who is not lying (sincere) and not mistaken (competent) will assert the truth (1994 p.147). This may not be uncontroversial, but it is useful for me to stick to this simple picture.

[2] As I said at the beginning of the thesis, I will not be considering scepticism as a serious alternative. However, some, such as John Locke, advocate such a view. Jonathan Barnes claims: 'No doubt we all do pick up beliefs in that second hand fashion, and I fear that we often suppose such scavengings yield knowledge. But that is only a sign of our colossal credulity: [it is] a rotten way of acquiring beliefs and it is no way at all of acquiring knowledge.' (from Welbourne 1994, p.297) Such a view is untenable, in my opinion.

how one can know what is asserted. The obstacle can be overcome by showing that there is a connection between what a speaker says, and its being true. This connection gives the hearer the epistemic right to rely on a speaker's word, it means that when we do form true beliefs based on testimony, they are justified[3] and can have the status of knowledge.

In this section I want to defend the claim that most approaches to the problem of testimony so far have been projects that fall within this broad picture. Their aim is to overcome the problem of vulnerability by presenting arguments that show that we are entitled, justified or warranted in our testimony-based beliefs. I will argue that a unifying feature of the approaches is how they see the goal of their project. Their project is to cross the epistemic gap produced by the indirectness of testimony; to establish an epistemic link between a person hearing a speaker assert that p and p's being true. This link, whatever it is, is what makes our testimony-based beliefs justified/warranted/eligible for the status of knowledge. The aim of their project is to overcome what appears to be a threat to the possibility of knowledge through testimony.

The different approaches to the problem of testimony are often divided into either two or three main camps. These camps are the reductionists, the non-reductionists, and the trust or assurance views that are sometimes classified as non-reductionist. I will deal with trust views separately as they represent a distinct way of looking at the problem of testimony.

The reductionists

Reductionism, whilst attracting many followers a decade or more ago, has somewhat fallen from grace since then. It is the view that the justification or entitlement that we get to accept testimony derives from other 'more basic' sources of knowledge. The idea is that knowledge via testimony gains its status as knowledge from perception, memory or inference (more 'direct' sources of knowledge). On the occasions when we can trust that

[3] I will use 'justified' or 'justification' in its broadest sense, from now on: a project which attempts to establish an epistemic link between hearing that p, and p's being true: a bridge to cross the epistemic gap – it would cover internal (e.g. reasons for belief) and external (e.g. reliability) types of justification. (Though it may exclude deontological conceptions of justification; but as far as I know, none of the epistemologists in the field of testimony take a deontological view of justification.)

what a speaker says is true, it is because we have evidence from our past experience, or from the evidence of our senses. This is what justifies our testimony-based beliefs.

The reductionist approach is an approach that is most clearly of the view that the way to overcome the obstacle to the possibility of testimony-based knowledge is to identify the justification that we have for believing that testimony is a good source of belief. The intuition that we need to know that testimony is a good source of belief before we can rely on it is a thought first expressed by David Hume. He argued: 'our assurance in any argument of this kind is derived from no other principle than our observation of the veracity of human testimony, and of the usual conformity of facts to the reports of witnesses.'[4] … 'The reason, why we place any credit in witnesses and historians, is not derived from any *connexion,* which we perceive *a priori*, between testimony and reality, but because we are accustomed to find a conformity between them.' [5]

The view that he defended is that we generally find testimony to be reliable, and from this we can infer that testimony is generally reliable, and thus a good source of belief. It is this inference from past experience that justifies our reliance on testimony. What has been pointed out, though, is that it is not really true that we have done any kind of checking on the general reliability of testimony, and the idea that we possibly *could* do such a thing has been brought into question. The Humean approach has been dismantled quite thoroughly by C.A.J. Coady in his book 'Testimony: A philosophical study.'[6] He argues that Hume's reductionist thesis is either 'circular' or 'simply false'.[7]

Jonathan Adler, in 'Belief's Own Ethics',[8] defends the claim that although we do not need to *check* on the reliability of testimony, we do have access to an 'enormous amount of background evidence'[9] that testimony is generally reliable. All of our experience of testimony, its general coherence, the fact that we are socially constrained to speak truthfully, the fact that we know that informants are by and large competent, etc., all this justifies our reliance on testimony as a good source of belief. So even though we do not, and could not, check on the reliability of testimony in general, we still have access to good

[4] ed. Millican 2007 pp.80-81.
[5] *Ibid*. p.82.
[6] 1992.
[7] *Ibid.* p.93.
[8] 2002.
[9] *Ibid.* p.153.

reasons to think that testimony is reliable. By this reasoning, Adler argues that our testimony-based beliefs are justified.

Elizabeth Fricker, one of the main modern proponents of reductionism, thinks that we cannot have empirical access to the reliability of testimony *in general*, but that this does not mean that we cannot find justification for thinking that *a particular instance* of testimony is reliable.[10] She argues that although we cannot check on the reliability of testimony in general, there are signs available in each instance of testimony that count as *evidence* that a particular speaker is trustworthy. She calls this 'local reductionism', as opposed to 'global reductionism' which aims to establish that testimony is generally reliable.[11]

What distinguishes the reductionists as a group is the view that hearing a speaker assert that p does not offer its own kind of justification. Whatever reason or justification we have for thinking that p is true comes from outside of the 'testimony event'. The justification we have for relying on testimony *reduces* to other more basic forms of justification, such as inference or perception. For Fricker: 'It can be the case that, on a particular occasion O when a speaker S makes an utterance U and in doing so asserts that P to a hearer H, H *has, or can gain, independent evidence* sufficient to warrant her in taking S to be trustworthy with respect to U.'[12] if the speaker is trustworthy, the hearer will be in a position to tell that she is trustworthy. She argues that although understanding an assertion is not sufficient grounds for coming to know what is asserted, the hearer has access to more than just the fact that the speaker has asserted that p: the fact that the speaker is trustworthy as regards p is an 'empirically ascertainable matter.'[13]

It seems quite clear that reductionism, in all its forms, falls into the category of attempting to overcome the obstacle of vulnerability. The presumption that the obstacle to knowledge posed by its second-hand nature can be overcome is explicit: there is more than just hearing an assertion that p that grounds a hearer's resulting knowledge that p. The idea behind reductionist approaches is that there is evidence to support our testimony-based beliefs. Even though our knowledge depends on a speaker getting it right, that is, not lying or being mistaken, we are justified (on the occasions that the justification is available) in

[10] 1994.
[11] *Ibid.* pp.133-134.
[12] *Ibid.* p.133 (emphasis added).
[13] *Ibid.* p.140.

accepting that what she says is true. The reductionists explicitly set out to identify our justification for taking a speaker at her word, and thereby show how knowledge through testimony is possible.

It can be seen that vulnerability is recognised as the problem facing testimony as a source of belief, and is seen to be a problem that can be surmounted by the presentation of reasons or arguments that support the epistemic status of testimony-based beliefs. It is accepted that the second-hand nature of testimony might count as a threat to the possibility of testimony-based knowledge, but if we establish that these beliefs are in fact justified, then knowledge is still possible. The justification that we get, either from our past experience, or the general coherence of testimonial reports, or signs of a particular speaker's trustworthiness, is what overcomes the obstacle.

The non-reductionists

The non-reductionists differ from the reductionists in that they believe the kind of justification or entitlement that we get from an instance of testimony is not reducible to other sources of knowledge. Our justification for accepting testimony does not depend on inference, perception or external evidence of any kind. Instead, testimony offers its own special kind of justification. There is something about testimony that it brings along with it a reason to take it to be a good source of belief. C.A.J. Coady was perhaps the first modern non-reductionist. He argues that there is a *necessary* connection between testimony and reliability, an *a priori* one. Coady argues for the thesis that testimony cannot be unreliable, because the very possibility of language use and communication depends upon the fact that testimony is mostly reliable. He claims that the reductionist approach to testimony rests on a presumption that it is possible that one's past experience or inference *could* discover that testimony is not reliable. He argues that this could not be so.[14] He argues for an anti-individualistic epistemology, that sees knowledge as something that is shared in linguistic societies.[15] Testimony is not a source of knowledge that is inferior to or dependent upon other more direct sources; it's being reliable is not something that we need to establish via

[14] Chapters 4-7.
[15] esp. Chapters 8 and 9.

other sources of knowledge. Testimony is a means of communication, and in order for it to function as such it is necessary that it is a reliable source of information.

Tyler Burge, another non-reductionist, argues for an *a priori* connection between intelligible testimony and truth. His thesis is that '*A person is entitled to accept as true something that is presented as true and that is intelligible to him, unless there are stronger reasons not to do so*.'[16] He does not claim, as Coady does, that testimony must necessarily be more often true than not. His view is not that testimony is reliable, as such, but rather that there is an *a priori* connection between an assertion's intelligibility and its being true. He argues that intelligibility is a sign of rationality, and rationality is, by nature, a source of truth.[17] This means that as long as we can understand an assertion (it is intelligible), then we are *a priori* entitled to accept it as true, because it comes from a rational source. We are entitled to rely on a rational source: just as we are entitled to rely on our own rationality as a source of truth, we are entitled to rely on the rationality of others as a source of truth.

The aim of the non-reductionists is to establish that we are entitled to our testimony-based beliefs based on reasons to do with the practice of testimony itself. The idea is that we don't need to resort to more basic or more 'direct' sources of knowledge to justify our acceptance of testimony. Unlike the reductionists, the non-reductionists hold that the indirectness of testimony does *not* mean that our warrant or entitlement to accept an assertion as true must be sourced elsewhere. Instead, testimony, by its very nature, is a source of truth.

The non-reductionists differ from the reductionists in an important way: they see our entitlement or justification for accepting testimony as being part and parcel of the practice of testimony, not external to it. This could be grounds for arguing that non-reductionists are dissipationists because they accept that all a hearer has access to is the assertion that p, and thus are arguing that understanding an assertion *is* sufficient for coming to know what is asserted. However, to argue this would be to misinterpret their claim. They do not straight forwardly claim that understanding the assertion is sufficient grounds for knowing what is asserted; they claim that the grounds for knowing what is asserted is provided by the a priori link between testimony and reliability, or between intelligibility and truth. Merely understanding an assertion does not give the hearer knowledge of what asserted, it is the

[16] 1993 p.467.
[17] *Ibid.* p.475.

fact that assertions are necessarily normally true, or that we are *entitled to presume* that an assertion is true, that can give our resulting belief the status of knowledge.

Non-reductionists do not differ from reductionists in how they see the problem of testimony: it is a problem of epistemic vulnerability. The way that non-reductionists go about responding to the problem of testimony demonstrates that they, like the reductionists, are committed to *overcoming* the problem. The non-reductionists are attempting to establish the link between hearing a speaker assert that p, and p's being true. They attempt to establish this link by showing that testimony is necessarily generally reliable, and thus a source of truth (Coady) or that intelligible assertions are by their very nature a sign of rationality, and rationality by *its* very nature is a source of truth (Burge). These connections are meant to provide the missing link between understanding an assertion and coming to know what is asserted. If our testimony-based beliefs are justified, or if we are *a priori* entitled to rely on others, then there is a link made between testimony and its being a good source of belief, and therefore a source of knowledge.

Non-reductionists, unlike reductionists, are not offering a story of justification *for the hearer*. It is not their aim to establish that a hearer is in a position to justify his testimony-based beliefs. The necessary reliability of testimony need not be something that the hearer knows about. The story of epistemic entitlement that the non-reductionists offer is theoretical, and not about the hearer's *reasons to believe* the speaker. But this does not mean that they are not attempting to overcome the problem of vulnerability. The goal of the non-reductionist is the same as the goal of the reductionist: he aims to *shore-up* or *strengthen* the epistemic status of testimony-based beliefs in the face of the threat posed by its second-hand nature. They are not offering first-person justification, but they are offering justification in the broad sense of providing a bridge to cross the epistemic gap between hearing that p and p's being true, a gap seen to be created by the second-hand nature of testimony.

Non-reductionists, like reductionists, are interested arguing that *despite the epistemic gap brought about by the indirectness of testimony*, knowledge can still be had. They are interested in overcoming the obstacle by showing how the insufficiency of mere understanding as grounds for knowing what is asserted is not incompatible with coming to know what is asserted. The entitlement provided by the a priori link between an assertion and its being true, or the necessary reliability of testimony, overcomes the problem of

epistemic vulnerability: knowledge can be had from testimony because of this justification (entitlement) we have for taking an assertion as true.

Trust and Assurance views

Trust and assurance views differ from the reductionist and non-reductionist approaches in that they reject the presumption that what is needed to overcome the obstacle is evidence or a reliability story. The idea is that the attempt to establish that testimony is a reliable source of belief is to miss the point of testimony, to misunderstand exactly what is involved in an instance of testimony.

The trust and assurance views hold that the reductionist and non-reductionist views are so intent on justifying the belief *in the proposition* offered that they miss that what is believed is not the proposition, but *the speaker*.[18] Looking at testimony as *evidence* for the truth of the proposition gets testimony wrong. We can get knowledge from testimony, not because a speaker's assertion that p is evidence that p is true, but because a speaker is giving us her assurance that p is true, or, as Paul Faulkner has it, a speaker telling us that p is giving us a reason to *trust* her.[19] The idea behind this break with what they call the 'evidentialist' approaches, is that the kind of reason we have for taking testimony at face value (for accepting it as true) is of a epistemologically distinct kind. It is not about evidence or reliability, it is about *interpersonal* reasons to believe. One of the reasons for thinking this is that if we consider what it is for a speaker to tell a hearer that p, it is explicitly *not* a case of the speaker giving the hearer evidence that p. The speaker is telling the hearer that p, and expecting the hearer to *take her word for it*.[20]

Richard Moran claims that if an assertion that p was *evidence* for the truth of p, 'telling' someone that p would be in the same category as showing him a photograph of p.[21] One can produce evidence that a certain state of affairs obtains, but it does not seem right to say that *saying* that it obtains is producing evidence that it does. It misdescribes the situation to see testimony in this light. The hearer believes the speaker based on the fact

[18] See Moran 2005, Ross 1986.
[19] Faulkner 2007.
[20] Moran 2005 p.6.
[21] *Ibid.* pp.10-12.

that the speaker takes a certain epistemic responsibility, and *assures* the hearer that what she says is true. It is this assurance that gives the hearer a reason to believe the speaker: 'On the Assurance View, dependence on someone's freely assuming responsibility for the truth of P, presenting himself as a kind of guarantor, provides me with a characteristic reason to believe, different in kind from anything provided by evidence alone.'[22]

Paul Faulkner modifies Moran's view, arguing that a speaker's assurance cannot count as a reason to think that what the speaker says is true, because assurance will be given *even when the speaker is lying or mistaken*.[23] A speaker's taking responsibility for the truth of what she asserts is not sufficient to give the hearer good reason to think that what she says is true. He argues instead that a hearer's trust in a speaker is warranted due to 'trust-based reasons.'[24] What he means by this is that affective (as opposed to predictive) trust in a speaker (trusting that the speaker will behave in a trustworthy manner) engenders that trustworthy behaviour. The hearer's epistemic dependence on the speaker gives the speaker a reason to be trustworthy. In the same way that a woman trusting an ex-con that she has hired not to steal from her *motivates* the ex-con to abstain from stealing, so a hearer's trust can motivate a speaker to make true utterances.[25]

Although it is true that the kind of approach that is taken here is distinct and interesting, I believe that it is still setting out to overcome the obstacle facing the possibility of testimony-based knowledge. The idea is that the speaker's assurance (Moran), or trust-based reasons (Faulkner), are what make the hearer entitled to accept that what the speaker says is true. There is something epistemologically distinct about 'the speaker's word' (it is not *evidence* that what she says is true), but the trust theorists are on the same path as the 'evidentialists'. The reasons they identify are not reasons about reliability and good evidence, but the reasons are still given to fill the gap between hearing an assertion and coming to believe it as true. The trust theorists aim to justify (in the broad sense) our testimony-based beliefs.

The problem of testimony as it is perceived by the trust and assurance theorists is that there is an epistemic gap between hearing a speaker assert that p and p's being true –

[22] *Ibid.* p.7.
[23] 2007 p.878.
[24] *Ibid*. p.879.
[25] *Ibid*. p.885.

understanding an assertion cannot, by itself, ground the resulting belief as knowledge. What they seek to do is to cross that gap by arguing that a speaker has special reasons to assert only what she thinks is true. There is a moral or affective side to testimony, such that these reasons can count as an epistemological basis for testimony-based beliefs. It is possible that a speaker may be lying or mistaken, but the fact that she has reasons to be trustworthy, or the fact that her assertion comes along with an assurance of trustworthiness, gives the hearer an entitlement or warrant to accept that what she says is true. The beliefs we have gained through testimony are *justified* because of the special assurance that a speaker offers along with her testimony, or because of the reasons that she has to be trustworthy. The epistemic gap brought about by the indirectness of testimony can be crossed by trust.

**

The accounts above have been given short shrift, I have not spent any time making them plausible or filling out important details. However, my aim here is not to oppose these views *per se*, my aim is to bring into question the goal of their project. I want to suggest that providing or establishing justification is not the way forward, not that their accounts do not provide or establish good justification. For this reason, the substantive content of their strategy is not relevant. I hope that what I have said is enough to suggest that what unifies these different approaches to the problem of testimony (reductionism, non-reductionism and trust/assurance theories) is that they all see the indirectness of testimony as posing a threat to the epistemic status of testimony-based beliefs – the problem to be solved is a problem of vulnerability. Each, in their own way, attempts to overcome the obstacle by showing how it is that our testimony-based beliefs are justified enough to count as knowledge. These arguments establish that, although understanding an assertion is not sufficient grounds, on its own, for coming to know what is asserted, there are other factors that explain how we can still get knowledge from the word of another. The possibility of testimony-based knowledge might have been threatened by the fact that it was second-hand, by the fact that it, necessarily, comes from another person who may lie or be mistaken. But that threat of vulnerability can be overcome if we see that our testimony-based beliefs are, to the contrary, well justified, or that we are epistemically entitled to rely on the word of another.

The justification offered by the different approaches vary in important ways. That they are all justifications in the broadest sense is discernible only in that the aim of their project is to show how the indirectness of testimony does not mean that the beliefs we gain thereby cannot count as knowledge. The project embarked upon by these epistemologists is a project of establishing the link between hearing testimony (understanding an assertion) and coming to accept it as true. It is presumed that the problem of testimony is a problem of epistemic vulnerability brought about by indirectness. If knowledge is to be had through testimony, it is assumed, we must be able to show how our testimony-based beliefs are justified, or why we are entitled to them, or how it is that the possibility of the speaker lying or being mistaken is not enough to undermine the possibility of getting knowledge in this way. The above approaches set out to overcome the obstacle to testimony-based knowledge by establishing that the insufficiency of merely understanding an assertion as grounds for knowledge is *not incompatible with* the possibility of getting knowledge in this way nevertheless.

Section Two: The limits of the overcoming response

It seems, from the above discussion, that the problem of testimony is generally understood to be a problem of epistemic vulnerability: understanding an assertion is insufficient grounds (or inadequate justification) for coming to know what is asserted. It is because the problem is understood as such that the demand placed on the epistemologist is one of providing justification.[26] The fact that testimony is indirect is seen to threaten the possibility of knowledge through testimony, and yet this threat is meant to be overcome by justification. Overcoming the obstacle is to deny that 'Merely understanding an assertion is insufficient grounds (inadequate justification) for coming to know what is asserted' is incompatible with the possibility of getting knowledge through testimony. The project embarked upon by most epistemologists in the field is a project of justifying our testimony-based beliefs, or establishing our epistemic entitlement to rely on the word of others and

[26] One notable exception is John McDowell's 'Knowledge by Hearsay', (1994), in which he argues that a hearer can just 'pick up' the knowledge that a speaker puts out there in a case of testimony, and he offers no justification, entitlement or reliability story. It is something like his story that I wish to defend; I endorse his claim that knowledge is there for the taking. But unlike McDowell, rather than just claiming that it is so, I intend to show how such a thing is possible.

thereby establishing the possibility of testimony-based knowledge. If the status of testimony-based beliefs as knowledge is seen to be at risk, the way to overcome this is to establish a justification, or show that testimony is, in fact, reliable, or that we are epistemically entitled to rely on the word of others. Establishing this, it is presumed, is the means by which we establish the possibility of testimony-based knowledge.

The overcoming approach is about accepting that there is a threat posed by the indirectness of testimony and yet working to overcome this by denying that this is incompatible with getting knowledge in this way. Knowledge through testimony is still possible despite the threat posed by indirectness because we are still warranted/entitled to the beliefs that we get through testimony. Yes, understanding an assertion does not, on its own, ground our knowledge of what is asserted, but there is something else that can explain how, despite this, testimony can give us knowledge. The project embarked upon by the above approaches to testimony is a project of denying that the epistemic vulnerability of testimony as a source of belief is an impediment to the possibility of getting knowledge through testimony. The presumption of such views is that establishing the justification or entitlement that we have to accept the word of another is *sufficient to establish that knowledge through testimony is possible*: the obstacle can be overcome.

It is this presumption that I wish to question in this section. One can establish the possibility of justified belief through provision of justification, I want to argue, but one cannot establish knowledge in the same way. The obstacle to testimony, that merely understanding an assertion is insufficient grounds for coming to know what is asserted, can be overcome if the justification provided *is* sufficient grounds for coming to know what is asserted. But knowing something entails that what one believes is *true*; being justified in one's belief does *not* entail that it is true. I want to suggest that this distinction causes a problem for those wishing to establish the possibility of testimony-based knowledge by attempting to overcome the problem via justification. The approaches discussed above provide a story that establishes justified belief, but this, I submit, is not the same as establishing the possibility of testimony-based knowledge. Overcoming the problem of apparent *unjustified belief* is not to fully overcome the problem of the threat to knowledge that (on their view) the indirectness of testimony poses.

The idea that understanding an assertion is not sufficient grounds (not justification enough) for coming to know what is asserted constitutes a threat to *both* justified belief and

to knowledge. Moran, for example, claims the problem is a clash between the 'classic empiricist picture of genuine knowledge basing itself either on direct experience of the facts or on working out conclusions for oneself' and testimony being 'so flimsy, unregulated, and known in plenty of cases to be unreliable, even deliberately so.'[27] The flimsy basis/grounds/justification for testimony based beliefs is seen to be what threatens the possibility of testimony-based knowledge. The presumption is that the threat to knowledge *is* the threat to justified belief posed by the indirectness of testimony. As a result it is thought that if one can provide a justification, or establish that the grounds we have are not weak or flimsy, then one can establish that knowledge is indeed possible.

However, even if it is true that a threat to the possibility of justified belief results in or entails a threat to the possibility of knowledge (which is what the dissipatory approach will deny), there is reason to doubt the conclusion that establishing justified belief is tantamount to establishing knowledge. This is because justified belief is not equivalent to knowledge. Under the most commonly accepted conception of knowledge, justified belief is a necessary but not sufficient condition for knowledge. Being male is a necessary condition on being a bachelor, and so not being male disqualifies one as counting as a bachelor. However, establishing that one is a male is not sufficient to establish that one is a bachelor. One also has to be unmarried. In the same way, being unjustified is a threat to the possibility of knowledge, but establishing a belief as justified does not thereby establish it as knowledge. The justificatory response to the problem of vulnerability establishes how we can get justified belief from testimony, but not how we can get knowledge.

The broad picture of knowledge within which the above epistemologists are working will probably be something like: X has knowledge that p if and only if X's belief that p is justified (broadly construed), and X's belief that p is true. The threat to the possibility of testimony-based knowledge, brought about by its necessary indirectness, is seen to be a threat to the first condition. The presumption is then that establishing that testimony is a justified source of belief (that we have evidence to accept it, that it is a reliable source of belief, that we are *a priori* entitled to accept it, etc.) then we have removed the threat to knowledge, and thus have shown how testimony-based knowledge is possible.

[27] 2005 p.1.

The implication of taking one of the approaches discussed above is that the establishing of the first condition as true allows for the argument to go through: knowledge is made possible. But in order for the justificatory response to be completely satisfactory in responding to the problem of how we get knowledge from testimony, it needs to be the case that the justification is the only thing threatened by the indirectness of testimony. If establishing justified belief is to be tantamount to establishing knowledge, establishing justified belief must be not only necessary but also *sufficient* for establishing knowledge: the threat to justified belief that the indirectness of testimony poses is the *only* threat to the possibility of knowledge. There needs to be a disanalogy with the bachelor case such that establishing the belief as justified is *all* that needs to be done in order to establish it as knowledge. And this, I contend, is not easy to do.

The second condition in the analysis of knowledge above is that the belief in question is *true*. If knowledge is understood to be justified true belief, then the belief must be true as well as justified in order to count as knowledge. One way to argue that establishing justification is tantamount to establishing knowledge is to argue that establishing justification establishes that a belief is true (that establishing the first condition is to establish the second as well). If the justification were able to establish that a belief was true, then establishing it as justified *would be* establishing it as true, and therefore as knowledge. But the responses to the problem of testimony, the approaches discussed above, cannot establish that what a speaker says is true. Elizabeth Fricker's local reductionist account can show us that we can have evidence supporting the trustworthiness of the speaker; Coady's theory about the necessary reliability of testimony can show us that it would be impossible for communication to exist if testimony were not reliable; Faulkner's trust based reasons are reasons we might have for thinking that a speaker will not lie or deceive us. But none of these can show us that what the speaker says *is true*. This is the same as saying that none of these can establish my belief in the testimony of another as *knowledge*. It is possible that some kinds of justifications may establish a belief as true, perhaps self-evident or a priori justification may be able to do this (e.g. the justification one has for believing that 2+2=4). However, none of the justifications offered in the epistemology of testimony have that kind of strength. The justificatory stories cannot fully explain how one can get knowledge (as opposed to justified belief) through testimony because establishing that the first condition holds cannot establish the second.

Another way to argue that establishing justification is tantamount to establishing knowledge is to argue that the indirectness of testimony poses a threat to the first condition, but *not the second*. In other words, one could argue the truth of a belief is not threatened by the possibility of a speaker lying or being mistaken, it is only the justification that is threatened. The obstacle to knowledge posed by the indirectness of testimony is that understanding an assertion is not *sufficient grounds* or adequate *justification* for coming to know what is asserted. The sceptical threat is to the justification condition only. Once an epistemologist establishes that this threat or challenge can be overcome, then the threat to knowledge is thereby removed. This would be analogous to claiming that the threat to a man being a bachelor is the question of his marital status. Once we have established that he is not married we have established that he is a bachelor. The question of him being a *man* was never raised.

But this would require that the truth condition for knowledge is not threatened by the indirectness of testimony. It would mean that the obstacle (merely understanding an assertion is insufficient grounds for coming to know what is asserted) is a threat to the justified status of testimony-based knowledge, but does not threaten the truth requirement. But why should we think this? The obstacle, that mere understanding is insufficient grounds for coming to know what is asserted, is an obstacle for a reason. It is an obstacle because *the possibility of a speaker lying or being mistaken is seen to undermine a belief's epistemic status*. Unlike when one uses one's own sense perception or inference, one's knowledge is *dependent* on the speaker *not lying or being mistaken.* The worry arises because knowing what the speaker has asserted is not enough to establish that what she asserts is *true*. A speaker could be lying, a speaker could be ignorant about what she claims to know: merely understanding what she says in insufficient grounds for knowing that what she asserts is true. The problem of vulnerability is a problem because understanding an assertion is insufficient grounds for establishing that *the truth condition* on knowledge has been met.

Perhaps the internalist would argue that the possibility of the speaker lying or being mistaken is a an obstacle to knowledge because it makes it *likely* that an assertion will be false, or because it makes testimony seem an *un*reliable source of knowledge. The problem is not that an assertion *could* be false, but that it is *probably* or *very possibly* false. This move would ensure that a justificatory response, the establishing of an epistemic link

between an assertion that p and p's being true, would be sufficient to overcome the obstacle. All that's lacking is the *probability* (or entitlement to accept etc.) that what is asserted is true, and once this is countered with provision of justification, the story about how we get knowledge from testimony is complete.

But the obstacle to getting knowledge from testimony is not that it is *likely* that an assertion will be false. The obstacle to getting knowledge from testimony is not a substantive claim about the general unreliability of testimony. The obstacle to getting knowledge from testimony arises *directly* out of the fact that it is second-hand knowledge, and its second-hand nature means that we must depend on a speaker not lying and not being mistaken in order for us to get knowledge of what she asserts. The fact that we are dependent on the speaker's competence and sincerity is the point at which the problem arises. The problem of vulnerability arises because the speaker *could* be lying, or *could* be mistaken, and thus merely understanding her assertion is insufficient grounds for coming to know what she asserts.

Seeing the problem of testimony as a problem of vulnerability is to presume that the epistemic status of a testimony-based belief, and (thereby) its claim to knowledge, is undermined or threatened by the possibility of things going wrong. The possibility of false testimony threatens the possibility of getting knowledge in this way. The fact that a speaker could be lying or mistaken directly threatens the truth requirement on knowledge. The problem of vulnerability is not just a problem for the justification condition, it is a problem for the justification condition *because* it is a problem for the truth condition. If there was no threat to the truth condition (if the possibility of false testimony did not threaten the possibility of testimony-based knowledge) then it is hard to see how the justification condition could be threatened. If it is not an epistemic threat to knowledge that the speaker could be lying or mistaken, then why should understanding an assertion be considered insufficient grounds for believing what is asserted? It is precisely *because of* the possibility of the speaker lying or being mistaken that the question of how we are justified arises (the threat to the second condition results in a threat to the first).

The indirectness or second-hand nature of testimony threatens the possibility of knowledge not only by threatening the justificatory status of testimony-based beliefs, but also by threatening their being *true*. The threat of vulnerability cannot be confined to being a threat to justified belief: the threat to justified belief is inextricably linked to the threat to

the truth of our testimony-based beliefs. One cannot claim that the problem of vulnerability threatens only the justification condition, that it undermines only the condition that we are entitled to our testimony-based beliefs, and does not undermine the truth condition. The fact that a speaker could be lying or mistaken threatens the grounds we have for believing her because it threatens the truth condition. If the indirectness of testimony threatens the truth condition of knowledge, and justification cannot establish a belief as true then establishing justified belief, as the approaches discussed above, does not establish the possibility of knowledge.

If the threat to knowledge posed by the obstacle is a threat to the truth condition *and* the justification condition for knowledge, then establishing a belief as justified will not fully or satisfactorily overcome the obstacle unless it can also overcome the threat to the truth condition. If the possibility of knowledge through testimony is threatened by the possibility of false testimony, then as long as the possibility of false testimony survives, so too will the threat. If understanding an assertion is insufficient grounds for coming to know what is asserted due to the possibility of the speaker lying or being mistaken, then supplementing this with having a priori entitlement, or the claim that testimony's is necessary reliability, will not be sufficient grounds for coming to know what is asserted either.

One could argue that, although the vulnerability problem threatens the second condition, this does not necessarily mean that one cannot establish the possibility of knowledge through justification. It is generally accepted amongst epistemologists that justified belief is not sufficient for knowledge: the second premise, that the belief be true, also needs to be met. But truth as a condition for knowledge is considered to be something *external* to the believer, something in the world. For this reason, it cannot be expected of an epistemologist to establish a belief as true. The most that one can do is establish justified belief, and justified beliefs are *candidates* for knowledge. In order for these beliefs to be knowledge they only require that the world cooperates. Once the world cooperates (the belief is true) and the believer has the required justification, then the belief counts as knowledge. In this way establishing justified belief is establishing the possibility of knowledge. Not all justified beliefs are true, but the ones that are, are knowledge. The epistemologist has done all that can be done by establishing the belief as justified, establishing the truth of a belief beyond providing justification does not even make sense.

What this route concedes is that one cannot establish beliefs as knowledge (as distinct from justified belief), but it is to claim that this is a necessary (and acceptable) consequence of the essential nature of knowledge: knowledge requires something not only of the believer, but something of the world. It is to claim that the problem of vulnerability is a only a threat to the justification condition because the truth condition is something external to the believer, and not something that needs to be (or can be) established. Once one overcomes the problem of a lack of justification, one successfully overcomes the problem of vulnerability. It cannot be the job of the epistemologist to establish beliefs as knowledge (as opposed to mere justified belief), and therefore an account which does not do so cannot be criticised for being unsatisfactory. Within the metaphysics of contemporary epistemology, there is no means by which one can establish a belief as knowledge (as opposed to justified belief) through justification, and therefore there is no means by which one can establish a belief as knowledge. One cannot demand the impossible. Justified belief is not sufficient for knowledge, but it is the most that one (as an epistemologist) needs to (and can) establish.

But this kind of concession is misleading in its simplicity. It is to claim that truth is a condition that can be established only by the cooperation of the external world, and therefore is something separable and distinct from the justification condition, and not, ultimately, within the purview of the epistemologist. It is to maintain that the role of justification can be conceptually distinct from the role of truth. But this is problematic. The reason that mere true beliefs are not considered to count as knowledge is because one can have a true belief accidentally, or by coincidence. If mere true belief was sufficient for knowledge, there would be no problem of vulnerability – the problem of testimony would not arise. Simply believing a speaker, and her assertion being true, would be enough for knowledge to be gained.[28] Understanding the assertion would be sufficient for coming to know what is asserted. But this is exactly what the justificatory approaches to testimony reject – it is their goal to justify our testimony-based beliefs precisely because the sufficiency of merely understanding an assertion is considered insufficient grounds for knowledge. It is presumed that just believing that p when one hears that p would not be

[28] It is worth noting, at this point, that the picture I defend later in this thesis rejects the idea that justification is necessary for knowledge. I argue that knowledge can be had from testimony without any epistemic entitlement/warrant/justification. But this is not to say mere true belief is sufficient for knowledge; knowledge is a mental state distinct from belief.

enough to gain knowledge of the external world because knowledge is meant to be a *special* or privileged kind of belief. Justification, as a condition of knowledge, is put forward as a way to overcome the vulnerability problem: knowledge is more than just true belief.

What this means is that justification is meant to *establish the link* between belief and truth, it is meant to provide the link between hearing an assertion that p, and p's being true. I believe that taking this to be the demand placed on one by the problem of testimony is incompatible with the concession that the epistemologist cannot establish a belief *as true*: it is to concede that whether or not one has knowledge is ultimately down to luck. A belief that is accidentally true is not knowledge, a lucky guess is not knowledge. But a belief that is justified, or grounded in the right way, or a belief that we are epistemically entitled to, is a belief that is true *not* by mere coincidence or luck, but *because* it is justified. The point of justification is to establish the link between a belief and its being true. If the fulfilment of the second condition (the truth condition) is just down to luck or coincidence, or the world doing us a favour, then how is a justified true belief any better that a mere true belief? Both, ultimately, must concede that whether we have knowledge or not is down to mere luck or coincidence.

What I am trying to make clear here is that the seemingly simple concession that the second condition (truth) is not something that can be established by the epistemologist, and that the goal of establishing the possibility of knowledge is completed once a justification is established, is not a simple concession at all. It is unsatisfactory, not only because the truth condition is not established, but because the establishing of the first condition (justification) is only a prima facie means to establishing the possibility of knowledge *because* it is meant to establish the link between belief and truth. Justification is posited as a condition on knowledge because it is unsatisfactory to imagine that a belief that is true by luck can count as knowledge. The fact that the justification does not eliminate the possibility of luck is inherently unsatisfactory; dismissing truth as a condition to be established is to dismiss, essentially, the point of establishing justification.

This same point can be made by following Gettier's lead, and showing how, intuitively, justified true belief is not sufficient for knowledge. It is always possible for a belief to be true and justified, and yet we would not count it as knowledge.[29] This is because the justification

[29] 1963.

has not established the belief as true. There are many possible responses to Gettier, but Linda Zagzebski has argued (successfully, I believe) that no modification to this model can succeed, and that 'Gettier problems are inescapable for virtually any analysis of knowledge which at least maintains that knowledge is true belief plus something else.'[30] Her argument succeeds because she shows that as long as it is always possible for the belief to be true *by luck*, then we will not wish to call it knowledge. The relationship between truth and justification is 'close but not inviolable',[31] and thus it is always possible for a belief to be true by luck. This is problematic because the point of justification is to distinguish between accidentally true belief and knowledge, and yet justification cannot show that a belief is not accidentally true. The above approaches to the problem of testimony cannot show how we get knowledge, as opposed to accidentally (though justified) true belief, through testimony.

Zagzebski's argument is that if a belief can be justified and yet false, then there is a possible example of a justified true belief, that happens to be true by accident.[32] So in the case of testimony, according to the reductionists, the hearer has evidence to the fact that a speaker is trustworthy or evidence that what she asserts is true. This justification is not enough to establish that what the speaker says is in fact true, so we can imagine a case where the hearer is justified but gains a false belief - the speaker is lying, for example. In order to 'Gettier' the example, we just need to make the belief true by accident. In this way, the hearer can have a justified true belief, and yet this belief would not, intuitively, count as knowledge. The speaker is lying but has got her facts wrong, and what she says is, in fact, true. The same is possible for the non-reductionist and trust approaches to the problem. The justifications they propose do not do enough to establish the possibility of knowledge: the second condition, the truth condition, could always be true by luck. The justificatory stories proposed as a response to the problem of vulnerability do not establish the possibility of knowledge – the justifications provided are *not sufficient* to cross the gap between understanding an assertion that p and coming to know that p.

But the idea that justification cannot establish knowledge (cannot establish that a belief is true) is not something new: 'any theory of knowledge that endorses the principle that S knows q on the basis of reason r only if r entails q, is doomed to a sceptical

[30] 1994 p.65.
[31] *Ibid.* p.65.
[32] *Ibid.* p.69.

conclusion. Fallibilist theories reject this entailment principle thereby avoiding this immediate sceptical result. The acceptance of fallibilism derives from the widely held view that what we seek when constructing a theory of knowledge is an account that squares with our strong intuition that we know many things.' [33] (This is taken from Stewart Cohen's paper, 'How to be a Fallibilist'.) The commitment to fallibilism,[34] the only logical response to the threat of scepticism entailed in the principle, is to concede that knowing something does not entail that one's justification eliminates the possibility of error.[35] This is something that almost all contemporary epistemologists have made (or have attempted to make) themselves comfortable with. Scepticism is untenable, and therefore we must accept fallibilism about justification. We must accept that justification cannot do the job that it was originally designated for, that is, to establish a belief as knowledge.

On the one hand this response just seems inevitable. But on the other hand there is something deeply unsatisfactory about it: it is to concede that the epistemologist cannot establish the possibility of knowledge. She can explain how it is that my testimony-based beliefs are justified, or why I am entitled to rely on the word of another, but she cannot explain how it is that I can *know* that p, when I hear it asserted. If knowing what the speaker tells me is something that occurs only when the world cooperates (when what she tells me is true), there is nothing to add beyond justification to *explain* how my getting knowledge occurs. The belief must be true, but this is a fact external to my being justified in believing it, and my being justified in believing it does not give me access to whether or not it is true. The justification I have for believing her is sufficient for justified belief, and a justified belief that is true is (supposedly) sufficient for knowledge. But the truth is something that may or may not coincide with the justified belief, and so the hearer knowing what the speaker tells him is something that may or may not occur when he believes what she says. There is nothing that the epistemologist can say that can explain how knowledge occurs as opposed to mere justified belief.

The fallibilist must make himself comfortable with the idea that knowledge is compatible with the possibility of error. Very few people would try to cling to the illusion

[33] Cohen 1988 p.91.
[34] Though I contend that calling it a *commitment* is misleading: it would seem more accurate to say one concedes to fallibilism, rather than commits to it. See Lewis 1996 for a similar view, p.550.
[35] Breuckner (2005), Dougherty and Rysiew (2009), Reed (2002), Stanley (2005).

that our capacity for reason is infallible, and few others would concede that we know nothing as a result. However, the point at which my argument is putting pressure is not just on the satisfactoriness of a belief's claim to knowledge through justification, but the satisfactoriness of a *project that explicitly sets out to establish the possibility of knowledge through justification*. It is inevitable that our reasons or justifications for a belief cannot establish a belief as true, but epistemologists of testimony set out to discover *how it is that we can get knowledge from testimony* (in the face of the possibility of false testimony) and yet cannot ever provide more than a justification for our testimony-based beliefs.

It is the *aim* of the project of overcoming the obstacle that contains the seeds of its own defeat. The aim is to establish knowledge, and this it cannot do. The aim is to provide the missing link between understanding an assertion that p and coming to know that p, but being justified in accepting p is not equivalent to knowing that p. Providing justification will not be enough to establish how we get knowledge of what is asserted. This is not because the justifications offered are not good enough, but because it cannot be the job of justification to establish a belief as knowledge. Justification is merely the reason one has to think that something is true, or the entitlement that one has to accept that something is true, or the epistemic warrant that one has to believe something: it does not and cannot entail that what is believed is true. Embarking upon a project of attempting to establish the possibility of knowledge through justification can never fully or completely succeed, because there will always be a gap between what counts as justified belief, and what counts as knowledge.

It seems to me that there is an inherent unsatisfactoriness in the justificatory response to the problem of testimony. The overcoming approach sees the obstacle as one of epistemic vulnerability and attempts to overcome this obstacle via provision of a justification. But this approach will only be satisfactory if the problem of vulnerability is *merely* a problem of lack of justification. But the problem is not merely a lack of justification: the possibility of knowledge is not threatened *only* because we have no justification for believing what we are told. The possibility of knowledge is threatened because in order to know something, it must be true, and the essential indirectness of testimony means that it is always possible that a speaker could be lying or mistaken (that what she says is *not* true.) The justification we have for accepting the word of another does not explain how we get *knowledge* from her.

At this point I can still feel the inclinations of the fallibilist pulling against me. The fallibilist has already conceded that knowledge cannot require the justification to entail the truth of the belief, so why must I continue to force an imaginary gap between establishing justified belief and establishing knowledge? Establishing justified belief *is* establishing knowledge because knowledge requires only 'probilification' or some kind of epistemic support: providing this is sufficient to establish the possibility of knowledge. But unless the fallibilist holds that there is *no* distinction between justified belief and knowledge, he still has a problem. If knowledge is something distinct from justified belief (which most fallibilists would concede – knowledge entails truth) then establishing justified belief does not establish knowledge. A belief has to be true in order to count as knowledge, and the threat of vulnerability facing testimony as a source of knowledge threatens the status of beliefs as knowledge: i.e. as true.

The problem of testimony - the problem of vulnerability posed by the indirectness of testimony - puts pressure on the fallibilist's ability to satisfactorily overcome the problem of vulnerability. This is because the epistemic threat to testimony-based beliefs arises out of the presumption that understanding an assertion is not sufficient grounds for coming to know what is asserted: if understanding the assertion is not enough, then *something else* must be able to cover that gap, and yet no reason or justification is such that it can establish a belief as knowledge. Seeing testimony as being faced with a problem of vulnerability exposes a commitment to the idea that we ought to or want to be able to establish how knowledge can be achieved. We want to be able to show that we can *know* that the speaker is telling the truth. We see this threat to knowledge and we want to show that knowledge can be had in spite of this threat. And yet embarking on said project cannot ultimately succeed because all we can ever show is that testimony-based beliefs are justified or that we are entitled to them.

In order for a response to the question of how testimony-based knowledge is possible to be satisfactory it must tell me *how it is that I know that p* upon hearing it asserted that p. The justificatory response to the problem does not tell me anything more than that I am entitled to accept testimony, that testimony is necessarily reliable, or that my testimony-based beliefs are warranted. But when I am told that Lima is the capital of Peru, or when my brother tells me that he hates his job, I *know* these things. It is not enough to be

justified in believing these things, a circumstance that would also be the case if what I was told was false. A satisfactory response to the question of how I get knowledge from testimony needs to explain how I get knowledge as opposed to mere justified belief. But the responses to the problem of testimony offered above can do no such thing.

This may be a pill that many are prepared to swallow, given the conception of knowledge with which they are working. No one expects justification to provide a belief with the status of knowledge because no justification is infallible, and in order to establish a belief as knowledge (as true), the justification must be (at least in the instance in question) infallible. But, whether or not one swallows the pill, it is still a pill. It is not a satisfactory or happy response to the problem. And the idea that it is inevitable is, I believe, a consequence of accepting the problematic internalist conception of knowledge. It is only if knowledge is presumed to be reducible to conceptually distinct internal and external components that this problem arises. If the truth is external from the knower, establishing knowledge in the face of a vulnerability threat is always going to be a frustrating task.

The epistemologist who sets out to establish the possibility of knowledge by establishing justification believes that: the lack of justification threatens the possibility of knowledge; and that the provision of justification can overcome that threat. However, although this position has a prima facie plausibility, it misses that accepting the first claim is to fall into a sceptic's trap. If one accepts that a certain means of gaining knowledge is not sufficient grounds for coming to have knowledge in this way, then one must accept that whatever justification they provide *must be sufficient* to establish how we get knowledge in this way: the justification must be able to establish how the beliefs gained in this way *count as knowledge*, and no justification has this kind of strength. Accepting that the problem of testimony is a problem of vulnerability is equivalent to accepting something like 'Seeing that p is insufficient grounds for knowing that p because one may be a brain in a vat.' Accepting this as a problem means that whatever grounds one provides to overcome it, they must be good enough to establish that we are not a brain in a vat. Accepting that there is a gap to be crossed by justification between having a belief that p and knowing that p places a demand on the justification to go *all the way* from believing it to knowing it. Such a task, I contend, is hopeless.

There is one last way to respond here, however, and that would be to argue that one does not *want* to establish knowledge, and that justified belief is, and should be, the *only* explicit aim of an epistemologist.[36] Whether or not one gets *knowledge* is of no epistemological importance. The threat of vulnerability, one could argue, is a threat to justified belief, and that in itself ought to delimit the goals of the project. Knowledge, as a distinct epistemic state to be established, should be taken off the table. If knowledge is not the explicit (or even implicit) aim of the epistemologist in responding to the threat of vulnerability, then the fact that the possibility of knowledge cannot be established is not a problem. Establishing justified belief in the face of the threat to justified belief would count as successfully overcoming said threat.

But what would warrant such a retreat? Perhaps one could argue that knowledge through testimony is impossible, and therefore one should not even attempt to establish it as being possible. This, I concede, would warrant such a retreat. However, this response amounts to scepticism about testimony-based knowledge. Such a view may be tenable to some, but, as I said in the very beginning of my thesis, I will not be taking such a view seriously. I do know when I was born, what planets constitute our solar system, who won the world cup (I didn't watch the final myself), etc. It is the aim of this thesis to explain how it is that this happens – how I get knowledge from the testimony of others - but it is not something that I doubt, or am willing to doubt, happens.

To allow that establishing knowledge is not possible is to allow that the threat to knowledge posed by the indirectness of testimony is one that cannot be overcome. And I do not think that most epistemologists in the field would be willing to concede this. Instead it is their aim to establish the possibility of knowledge *via* the establishment of justified belief. They do intend for the provision of justification/entitlement/ warrant to be a satisfactory response to the problem of vulnerability posed by the indirectness of testimony. Implicit in such an aim is the idea that establishing justified belief is *sufficient* to establish knowledge, and that a fallibilist conception of knowledge allows this to be the case.

[36] Bertrand Russell (ed. 1998) and Crispin Wright (1991) both argue for a retreat from the quest for knowledge to the quest for justified belief. Their response to sceptical challenges is to concede that we ought to be satisfied with warranted belief, or 'probable opinion.' (Russell, p.81 [1998 edition]). Wright says 'We can live with the concession that we do not, strictly, *know* some of the things we believed ourselves to know, provided we can retain the thought that we are fully justified in accepting them,' (1991 p.88). Littlejohn (2013) argues that this retreat is a 'failure to appreciate the normative significance of knowledge' (p.2) and I tend to agree with him on this.

But, contra this, I have argued that justified belief is not sufficient for knowledge, and nor should it appear to be so to anyone who sees the problem of testimony to be a problem of epistemic vulnerability. There is an innate incoherence in taking the problem to be one of vulnerability, and presuming that justification can satisfactorily overcome it. If the possibility that a speaker could be lying or mistaken is seen to threaten the epistemic status of the beliefs we get through testimony, then as long as that possibility obtains, so too does the threat. The threat to knowledge that the problem of vulnerability poses is not merely a threat to the possibility of justified belief (and therefore to knowledge), it is a threat to knowledge directly because in order to know something what is believed must be true, and what a speaker asserts may not be true.

My claim has been that the justification condition will never be enough to establish the truth condition, and that this is fatal to one's project if one wishes to use the justification to establish the possibility of knowledge. The familiar problem of the limited power of justification surfaces in many different guises, the problem of scepticism[37] and Gettier's problem being two of them. The conception of knowledge that gives rise to the problem (the idea that knowledge is reducible to true belief and some other condition) is not a comfortable resting place. There are, of course, many different ways of dealing with the problem, many different ways of cashing out the concepts involved, but none that are widely accepted.[38] There may still be a satisfactory escape from the problem from within the accepted paradigm, but the stickiness of the problem, and its constant recurrence is enough, in my book, to take seriously some alternatives to the accepted paradigm.

[37] See Stroud (2000) for arguments that scepticism is a persistent and persisting problem for epistemologists.
[38] I haven't dealt directly with a causal reliabilist response. One could argue that justification *does* entail the truth of the belief: if the justification is causal, for example, then there is a *necessary* connection between the belief and its being true (e.g. Alvin Goldman, 1967). But if this is the case then it is not just the truth of the belief, but whether or not it is justified that becomes external to the knower, and this is a strange picture of knowledge. One could reject the claim that the causal relationship (the necessary link to truth) is external to the knower (i.e. non-mental), and this is route taken by externalism about knowledge (the view that I endorse in the next chapter).

Conclusion

The aim of this chapter was to suggest that most approaches to the problem of testimony are attempts to overcome the obstacle to the possibility of testimony-based knowledge (the vulnerability problem) through justification (broadly conceived), and that these approaches cannot completely and satisfactorily achieve this goal. A common presumption that I have questioned is that establishing justified belief is tantamount to, or equivalent to, establishing knowledge. Often the project embarked upon by epistemologists in this field is explicitly a project of providing justification for testimony-based beliefs. What I have tried to argue in this chapter is that such a project is *not* equivalent to establishing the possibility of knowledge. Furthermore, accepting that the problem is one of vulnerability is to accept that the possibility of error poses a threat to knowledge. What I have attempted to reveal here is the internal incoherence that this position entails. Taking the route of justification to establish the possibility of testimony-based knowledge is, it seems to me, doomed from the start.

I believe that this incoherence arises out of the innate incoherence of the internalist conception of knowledge. One wants to establish the possibility of knowledge, but due to the nature of knowledge as a state comprising internal and external components, one is (as an epistemologist) only in a position to establish the possibility of the internal components of knowledge, i.e. justified belief – not knowledge *in toto*. The claim that justified true belief is a problematic conception of knowledge is a familiar one. However, something like this paradigm has been tacitly accepted by the epistemologists I have discussed in this chapter. The accusation that I have laid against their approach to testimony arises out of the problems facing this broad epistemological model. The problem of vulnerability facing knowledge through testimony is one that is faced by *all* knowledge. What the general unsatisfactoriness of the approaches to the problem of testimony suggest is that what is needed is a rejection of that broad epistemological model, a rejection of the idea that knowledge is reducible to something like justified true belief.

In this chapter, my intent was to defend the idea that the justificatory (overcoming) response is *not sufficient* to establish knowledge. In the next chapter I will argue that such an approach is *not necessary*. In the next chapter I will present, and attempt to make plausible, a view of knowledge that entails a rejection of the epistemological framework

that has proved so problematic in the last few decades. This conception of knowledge (as espoused by Timothy Williamson,[39] and John McDowell[40]) will allow us, I contend, to accept that testimony is necessarily indirect, and yet not accept that this inevitably results in epistemic vulnerability. Under Cassam's structured approach to how-possible questions, this will amount to a *dissipation* of the problem of vulnerability. If we understand knowledge in the right way, the possibility of error does not pose a threat to the possibility of knowledge. The fact that the knowledge we gain from testimony is derived from the knowledge of another person does not mean that the epistemic status of this knowledge is threatened by the possibility of that person lying or being mistaken. Knowledge is a factive mental state, and not undermined by the possibility of error.

[39] 1995, 2000.
[40] 1995.

Chapter Three: Dissipating the obstacle

The second stage of Cassam's structured response to how-possible questions is the stage of obstacle removal. The perceived obstacle to the possibility of knowledge is what gives rise to the how-possible question, and it must be overcome or dissipated in order to show how knowledge is possible. The perceived obstacle to the possibility of testimony-based knowledge is one of epistemic vulnerability: merely understanding an assertion is not sufficient grounds for coming to know the proposition asserted. Vulnerability is seen to be a consequence of testimony's essential second-hand nature. In order to get knowledge from a speaker, what she asserts must be true: it cannot be the case that she is lying or mistaken. The possibility of the speaker lying or being mistaken undermines or threatens the epistemic status of testimony-based beliefs *qua* knowledge.

In the previous chapter I argued that most approaches to the problem of testimony so far have been about *overcoming* that obstacle. The aim of these approaches has been to establish a connection between hearing a speaker assert that p, and p's being true. In other words, they seek to justify testimony-based beliefs, or establish that we are entitled to accept or warranted in accepting that what a speaker says is true. In this way, they attempt to establish how knowledge through testimony is possible. I also argued that there is reason to think that these approaches will be ultimately unsatisfactory in overcoming the obstacle of vulnerability in this way. Once it is accepted that the possibility of false testimony threatens the possibility of knowledge, this obstacle cannot be completely or satisfactorily overcome without getting rid of or eradicating the possibility of false testimony, and this cannot be done. Establishing justified belief is not equivalent to establishing knowledge – provision of justification is not *sufficient* to overcome the obstacle.

In this chapter I seek to defend a different response to the problem of vulnerability, and argue that provision of justification is not *necessary* to establish the possibility of knowledge. Instead of recognising the obstacle of vulnerability, and attempting to overcome it, I will deny that it exists. I will deny that the possibility of speaker error poses a threat to the possibility of testimony-based knowledge. This way of dealing with an obstacle, under Cassam's picture, is called dissipation. If the problem of vulnerability arises out of the possibility of error, then this problem is not a problem that is specific to testimony-based

knowledge. Almost all sources of knowledge are fallible in some way[1], and this means that the possibility of error, if seen to threaten knowledge, would threaten almost all of our knowledge. The way to respond to this is to show how the possibility of error does not need to threaten the possibility of knowledge, and the way to do this is, I will argue, by espousing an externalist conception of knowledge.

I agree that there is something epistemologically distinct and special about how we get knowledge through testimony, but what makes it special is not that certain things can go wrong. This is not special at all. The possibility of error is not specific to testimony, and reducing the problem of testimony to the problem of vulnerability is to miss what it is that makes testimony special. What makes it special is that, through testimony, knowledge is passed on from one person to another simply through linguistic communication: merely understanding an assertion *can explain* how one can come to know what is asserted. How this happens is the subject of chapters four and five. The subject of this chapter is the dissipation of the problem of vulnerability. In order to dissipate the problem of vulnerability I need to show how we can reconcile ourselves to the idea that testimony-based knowledge is not threatened by the possibility of false testimony.

In the first section of this chapter I want to make the externalist conception of knowledge more palatable.[2] It is a view that is still controversial, and not widely accepted within epistemology, to say the least. However, embracing this conception is a necessary step in my thesis, and though I do not hope to convince readers to embrace it, I do hope to make the view a plausible alternative to the more commonly accepted internalist conception. Once I have done this, I will attempt to show, in Section 2, how taking on this view of knowledge can allow one to dissipate or ignore the problem of vulnerability. If knowledge is not seen to be a composite state of internal and external components, but rather a mental state which encompasses the external, then I believe that the indirectness of testimony does not need to be seen to undermine or threaten the epistemic status of testimony-based knowledge.

[1] Perhaps all are, but this would be a contentious claim. We know, at least, that our senses and our memory can fail us at times.
[2] For a thorough defence of externalism, see Timothy Williamson's *Knowledge and its Limits* (2000).

Section One: The externalist conception of knowledge

The externalist conception of knowledge as I present it here is mainly attributable to Timothy Williamson. His book *Knowledge and its Limits*[3] is a treatise devoted to the explication and defence of such a conception of knowledge. However, he is not the only philosopher to espouse this view: John McDowell, who rejects the 'interiorization of the space of reasons'[4] is a proponent of a similar view. Jennifer Nagel also defends the idea that knowledge is an irreducible mental state[5], and Michael Welbourne, in his examination of the phenomenon of knowledge transmission through testimony, concludes that knowledge is not a species of belief as is usually thought.[6] There are also philosophers of mind, such as Gregory McCulloch, Colin McGinn and Mark Rowlands, who reject the metaphysical view of mind that separates the internal from the external, and argue that the mind should be understood to encompass, or be partly constituted by, the external world.

These philosophers have shaped my view of knowledge, and in what follows I will attempt to do their views justice. However, it should be noted that externalism as I present it here is simplified down to the bare essentials, and the richness and diversity within externalism will not be represented at all. Consequently, the particulars of Timothy Williamson's conception of knowledge as distinct from John McDowell's, for example, will not feature here. It may also be true that proponents of externalism referred to above will not agree with everything I say here. The main aim in my explication will be to distinguish it from internalism in a way that makes clear how it could dissipate the problem of vulnerability facing testimony. Perhaps, and I say this tentatively, my views in this regard coincide most closely with John McDowell's.[7]

The key shift demanded by externalism about knowledge is the rejection of the idea that knowledge reduces to mental and non-mental components: this involves rejecting the mental/non-mental (or internal/external) ontological divide. However, although this sounds radical, it is something that has been embraced by most philosophers of language for some

[3] 2000.
[4] 'Knowledge and the Internal' 1995, 'Knowledge and the Internal Revisited' 2002, 'Criteria, Defeasibility and Knowledge' 1982.
[5] 'Knowledge as a Mental State' 2013.
[6] *Knowledge* 2001, 'The Transmission of Knowledge' 1979, 'The Community of Knowledge' 1981.
[7] Especially 'Knowledge by Hearsay' 1994.

time. Content externalism, the view that '"Meanings" just ain't in the head!',[8] is widespread in philosophy of language. The idea is that content, or meaning, is not something that can be made sense of in purely psychological terms. We need the world to be a certain way in order for the words we use to have the meanings that they do. As Gregory McCulloch puts it: 'content or meaning is environmentally constrained: the meanings you grasp at least sometimes depend not just on how things are inside you but also on what your surroundings are like.'[9]

In his paper "The Meaning of 'Meaning'"[10] Hilary Putman argues for this with a vivid thought experiment. Imagine there is a parallel universe which is entirely identical to this universe, except for one small detail. In the parallel universe, on what Putman calls 'Twin Earth', the molecular structure of water is not H_2O, but XYZ. On *this* earth, when I think about water, I am thinking about water H_2O, but on Twin Earth, when I (the other me) think about water, I am thinking about water XYZ. From the inside, from the first person perspective, our thoughts are identical, and yet we are thinking about *different things*. We are both thinking about a clear liquid that comes out of the tap, that quenches thirst, and that covers a large proportion of the earth's surface. However the *content of our thoughts differ*. Water H_2O is not water XYZ. On one earth I am thinking about water, and, on the other, I am thinking about 'twater'.

If it is possible for what one is thinking about to be not completely transparent to the self, and at least partly constituted by the way the world is, then there is a blurring of the ontological distinction between the internal/psychological, and the external world.[11] Content externalism has been used to combat global scepticism, in arguments such as 'If I am a brain-in-a-vat, then my word 'tree' does not refer to trees, my word 'tree' refers to trees. So, I am not a brain-in-a-vat.'[12] However, the success of the arguments as anti-sceptical is not my interest here. What is important is the ontological picture that is being painted. It entails a rejection of a neat Cartesian divide between the internal (mental) and external (non-mental/world). The internal, or psychological, realm is not conceptually

[8] Putnam 1975, p.13.
[9] McCulloch 2003, p.41.
[10] 1975.
[11] McCulloch's *The Life of the Mind* (2002) – a book deeply influenced by McDowell – constitutes a rejection of this distinction.
[12] From Stanford Encyclopaedia of Philosophy, 'Skepticism and Content Externalism', Brueckner 2012.

independent of the external world. More than this, the internal realm *is partly constituted by* the external. The actual content of my thoughts depends on the world being a certain way.

This idea is not very controversial in philosophy of language, and yet, when it comes to epistemology, the claim that my knowing something means that my *mental* state is partly constituted by the world being the way I take it to be, remains highly contentious. The view of knowledge that still has the most followers is a view that separates what is true (in the world) from what is mental (in the head). I am not sure why there is this discrepancy between philosophers of language and epistemologists, but perhaps the idea that the meaning of words, or the content of thought, being dependent on the world in some way is acceptable, but to say that my mental state of knowing that p actually incorporates or encompasses *p obtaining in the world* is too radical a move. Content externalists do not say that the mental content of my thinking about trees entails the existence of trees – otherwise thoughts about unicorns, for example, would depend upon the existence of unicorns – but knowledge externalists do hold that my mental state of knowing that there are trees over there entails that there are trees over there *in actual fact*. The mental state entails the world being a certain way.

It seems to me that content externalism should pave the way for externalism about knowledge.[13] However, this seems not to have been the case. Perhaps another one of the reasons for the perceived 'radical-ness' of externalism stems from a preconception of what qualifies as mental. There is a popular idea that what is mental must supervene exactly onto the physiological self.[14] The 'narrow' conception of the mental has it that there can be no difference in mental state if there is no difference in physiological state.[15] However, the way that we talk about mental states in ordinary language belies such a conception of the mental. Mental states such as 'recognising', 'remembering that', 'seeing that', 'understanding that', cannot be understood in narrow terms. They are mental states that

[13] See Williamson 2004, where he argues against Sosa's view that content externalism does not result in an externalist conception of mind.
[14] Jaegwon Kim is famous for attempting to make sense of the mind by exploring the notion of 'supervenience': a notion that defines what is mental as *narrow*. The mental supervenes on the physical (the brain) in that any mental state has a physical base, and every mental change entails a physical change. The idea is that everything mental *supervenes* on the brain. (This is a small improvement on identity theory which holds that everything mental *just is* something physical.) See *Mind in a Physical World* (1999), and *Supervenience and Mind* (1993).
[15] Williamson 2000 p.52.

cannot be understood without reference to some state of the world. To understand that something is the case, or to see that it is the case entails that it is the case. The same goes for loving, hating, preferring and other relational verbs.[16] These are what Williamson calls 'broad conditions', they are mental and yet do not supervene only on the physiological self, but also depend on factors outside of the individual.[17]

The view that externalism espouses that internalism does not is that the mind is not just in the head.[18] It is to deny the metaphysical possibility of Descartes' evil demon, or the 'brain-in-a-vat' hypothesis. The mind is not something that can be separated from the external world, it is partly constituted by it. Descartes claimed: 'I am now seeing light, hearing a noise, feeling heat. But I am asleep, so all this is [could be] false. Yet I certainly seem to see, to hear, and to be warmed. This cannot be false; what is called 'having sensory perceptions' is strictly just this, and in this restricted sense of the term it is simply thinking.'[19] But, according to externalism, it is simply not the case that the mind only has access to 'appearances' or 'seemings': Descartes' retreat is not permissible. The seeing of the light, the feeling of the heat entails the existence of the light and the heat. Our experience of the world encompasses the world being the way we experience it to be.

Despite the majority of epistemologists rejecting such a view, this broad notion of the mental is not something foreign to our natural understanding of ourselves. In fact, I contend, it sits more easily with our notions of how we interact with the world than an ontological dualism that separates the mind from the world.[20] A broad notion of knowledge such that it can be understood as being at once purely mental, and yet not to the exclusion of the world, is a notion that sits well with our experience of the world as conscious agents within it. Externalism makes the world directly available to us. We no longer have to work from appearances and seemings and reasons to get to knowledge: it is just there, a part of our mental life. There is no divide between the mind and the material world: the mind encompasses the material world.

---

[16] 'that one loves Mary and that one hates Mary are broad conditions, for they depend on a relation to the particular individual named.' *Ibid.* p.53.

[17] In Chapter 2 of *Knowledge and its Limits*, he defends broad conditions as being purely mental, arguing that the internalist 'tries to find a narrow non-factive attitude that exhausts the mental reality underlying the broad factive attitude' but that 'they all prove to be inadequate' (p.54).

[18] McCulloch's thesis: 'the mind just ain't in the head' (2003 p.12) – a nod to Putnam's '"Meanings" just ain't in the head' (1975 p.13).

[19] 2008 p.19 (ed. Cottingham).

[20] And certainly better than monism.

The internalist view of the mental is that it is a realm of representations distinct from and yet based (somehow – mysteriously) upon the real world. Knowing that p is a coincidence of believing that p, which is a mental state, and p obtaining in the real world, which is not part of the mental state. There is a requirement that the external *coincides* with the internal in order for knowledge to occur. The two are realms are separable. When I see that there is a tree in front of me, there is some kind of 'picture' or mental representation of the tree in my mind, and there is a tree out in the world. There is also some kind of link between that tree being out there, and there being a picture of a tree in my mind, but this link, whatever it is, allows for the possibility for me to have this picture of the tree in my mind without there actually being a tree. Such a view of the mental, though commonly accepted by philosophers, seems quite removed from our actual experience of seeing a tree.

The externalist view, though philosophically controversial, is the idea that seeing the tree entails the tree being out there in the world. Knowing that there is a tree is not a matter of there being some match or coincidence between the tree out there and the mental image of the tree in one's mind. Instead the tree forms a part of the mental state of knowing that there is a tree. The actual tree is not excluded from the experience of seeing the tree, it is a necessary part of seeing the tree. The mental realm is not just what is in the head, there is no substantive divide between the mental and non-mental. The world is there, in our experience of it, just as we experience it to be. A conception of knowledge and the mind that forces us to retreat from the world, to become essentially distinct from it, is a conception that we should be suspicious of. Externalism, I contend, is much more in keeping with how we experience the world.

Of course there is a price to pay, and for most, it seems, this price is just too high. This view contradicts some of our most closely held philosophical intuitions, for example, when my mental state changes, my psychological state ought to change too. (By psychological state, I mean the purely transparent internal state – what it 'feels like' from the inside) This is not true if one is an externalist. If a person knows that he owns two cars, he will fail to know this if one of his cars gets stolen, but his psychological state will not change until he discovers this. This means that, for the externalist, his mental state can change, without his internal psychological state changing. Relatedly, it also seems wrong to admit that we do not always know what mental state we are in. If knowing is a mental state, we ought always

to know whether or not we are in it. Externalism does not allow this – sometimes one can think that one knows something (is in the mental state of knowing it) and yet not. Thirdly, (also relatedly) one must allow that luck enters what McDowell calls the 'space of reasons' (the mental realm, so to speak). When one fails to know something because the world is unluckily not cooperating, this is a failure within the 'space of reasons' not in the world. So, for example, when the owner of two cars unknowingly becomes the owner of one car due to one of his cars being stolen, his failure to know it occurs within his 'space of reasons', not just in the world.

The reluctance to let go of these intuitions will be strong, but it seems to me that there are times when we ought to let go of them. The first intuition is that 'if mental state changes, psychological state changes too.' But this is something that is denied by content externalism (which is widely accepted). The meaning of a word is dependent on the world being a certain way. When I am thinking of water I am in a psychologically identical state to my twin in a parallel universe who is thinking about twater: but we are thinking about different things. The same is true for any factive mental states, for example recognition. I recognise that person as Bob, the mental state of recognition depends on it being the case that the person I recognise is Bob. If it were not Bob (but his identical twin Bill) then this means I would not be in the mental state of recognising him, though I would be in the same psychological state. It doesn't always seem intuitively unacceptable that we can be in different mental states but the same psychological state.

The second intuition is that we ought always to know what mental state we are in. But Williamson shows this to be false in his argument against the luminosity of mental states.[21] He argues that if mental states were luminous then one would *always* be in a position to know if one were in a certain mental state, and this he shows is not the case. Even in the case of being in pain, there will be times when the difference between being in pain and not being in pain will be indiscernible, and one will not be in a position to know whether one is in pain or not. One can even think she believes something, and not actually believe it – such as a belief in God, or the trustworthiness of one's husband. One may think that she believes in God, but when this belief is tested it may become clear that she does not. One can think

[21] Chapter 4.

that one's intentions are altruistic when they are in fact selfish. The intuition that our minds are always transparent to us is not non-negotiable.

The third intuition is that we ought not to allow luck to enter the space of reasons.[22] Failing to know something when I have done nothing epistemically wrong is the fault of the world, the thought goes, and as such it is nothing to do with me or my mental state. But this intuition also ought to be rejected. One can think that one knows something, and not know it, and this is an *epistemic failure*, not just a failure of the world to cooperate. The unluckiness occurs *within* the mental realm, not outside of it. That we can sometimes think that we know something, and yet not know it is an acceptable feature of everyday life. What seems unacceptable is the idea that failing to know is a failure within the 'space of reasons' or within the mental realm (rather than a failure of the world). It is presumed that if we have done our epistemic best, then the failure to have knowledge is unlucky because *the world* has not cooperated. If I have done nothing wrong and believe things carefully and rationally, how can it be 'my fault' or a failure of my reason that I sometimes believe false things? Under the externalist conception, however, the view is that the failure does occur within the space reasons. Doing one's epistemic best is not always good enough, if one does not have knowledge then the *believer has failed* to grasp the facts of the matter. It is unlucky not because the world has failed to cooperate, but because one's faculty of reason has failed to do its job and ensure factiveness.

I believe that there is a parallel here between epistemic luck and moral luck, as Thomas Nagel discusses it.[23] Nagel shows that our moral judgements of actions depend not only on the internal state of the agent (on factors within his control – if there is such a thing), but also on factors beyond his control. It is the norm to take into account the *actual consequences* of one's actions when judging an action as right or wrong. So even though there is a sense in which one could not have done otherwise, one has still *done wrong* by having a false belief. One is still at fault when one thinks one knows that p, and yet p happens to be false. One can rationally believe something false, but one is still doing something epistemically wrong, even if it is unlucky that one's belief is false.

[22] McDowell's phraseology (1994 p.217).
[23] 1979: 'The things for which people are morally judged are determined in more ways than we at first realise by what is beyond their control.' (p.26).

One might not want to accept that when something goes wrong in the world this can count as an epistemic failure, but the flip side of the coin is that, if we do, then when nothing goes wrong, this counts as an epistemic success. If one is an *internalist* about knowledge, then knowledge is no more of an epistemic achievement than justified belief. And yet, as virtue epistemologists like to maintain, there is sense in which having knowledge is an achievement. As John Greco puts it: 'To say that someone knows is to say that his believing the truth can be credited to him. It is to say that the person got things right due to his own abilities, efforts and actions, rather than due to dumb luck, or blind chance, or something else.'[24] This is much more in keeping with our intuitions about knowledge than the idea that when we know something it is down to (essentially) a coincidence between mind and world. The externalists can make sense of this because knowledge is considered to be a purely mental state, not a state composed of internal and external components. The idea of knowledge as an epistemic achievement is harder to make sense of if whether one knows something or not (a belief's being true or not) ultimately depends only on the world, not the believer.[25]

Although externalism about knowledge seems to contradict some of our closely held intuitions, I believe that it is more in harmony with a lot of our intuitions about knowledge than the idea that knowledge is a hybrid of external and internal components. If we think simply about our phenomenological experience of the world, it is difficult to marry this with a conception of knowledge that relegates truth to something irretrievably external to us – a view that holds that the fact of the matter is something substantively distinct from our experience of it. Existential phenomenologists focus on just this question. How do we actually *experience* our relationship to the world? Jean-Paul Sartre writes of internalism:

> we have all believed that the spidery mind trapped things in its web, covered them with a white spit and slowly swallowed them, reducing them to its own substance. What is a table, a rock, a house? A certain assemblage "contents of consciousness", a class of such contents. ... The simplest and plainest among us vainly looked for something solid, something not just mental, but would encounter everywhere only a soft and very genteel mist: themselves. [26]

[24] 2004 p.111.
[25] Even Greco, who argues that epistemic virtue and knowledge are intimately connected, must allow that 'conforming to correct epistemic norms' (1990 p.265) does not *always* result in knowledge, and therefore whether or not one's virtuous belief is actually knowledge is ultimately down to the cooperation of the world.
[26] 1970 pp.4.

What this quotation highlights is the strangeness of trying to reduce our phenomenological experience of the world to the immaterial mental. Under this view, all *things* become immaterial 'contents of consciousness'.

Understanding knowledge as a hybrid of internal and external components means that there is nothing mental that distinguishes a belief that has all the right internal components for knowledge but is false, to one that is true. There is nothing *about the knower* that distinguishes her from the mere believer. We are in the same mental state whether what we believe is true or false. There is nothing about our experience of the world that encompasses the world as it is. And if this is true (which for all proponents of internalism, it must be) then there is a sense in which we really *could* just be a brain in a vat or living in the matrix. The world is not a necessary feature of our experience of it. The world becomes, in some fundamental sense, inaccessible.

What seems more phenomenologically accurate is:

> You see this tree, to be sure. But you see it just where it is: at the side of the road, in the midst of the dust, alone and writhing in the heat, eight miles from the Mediterranean coast. It could not enter your consciousness, for it is not of the same nature as consciousness. ... To know is to "burst toward", to tear oneself out of the moist gastric intimacy, veering out there beyond oneself, out there near the tree ...[27]

Although we do not all think in this kind of language, it seems right that we experience the world as out there, not as the content of our consciousness. Seeing a tree is being in a certain epistemic or perceptual relationship with the *actual* tree, it is not a mental appearance or 'seeming'. This view is more in keeping with an externalist conception of knowledge because what it allows for is for the world to partly constitute the mental.

The point I am trying to make with this digression into the phenomenological, is that the view that knowledge is a hybrid of external and internal is not a naturally attractive one. It does not sit well with how we experience our relationship with the world. It is unnatural to think of our everyday life, our interactions with the world, as mental experiences that are conceptually or substantively distinct from the way the external world actually is. Such a view is forced upon us by our acceptance of a certain prevalent metaphysics: the idea that the world is external to the mind. McDowell, in his paper 'Knowledge and the Internal',

[27] 1970 p.4.

claims 'The apparent plausibility [of the internalist account] is not intrinsic to the position, but reflects an assumed framework; when one looks at the position on its own, the plausibility crumbles away,'[28] and I agree with him.

Embracing an externalist account of knowledge also allows us to hold on to the intuition that when we know something, we cannot be wrong about it. When a person knows something, this does not just mean that she has done her epistemic best, and, all things going well, what she believes is true. Instead, her *mental* state of knowing it ensures that it is true. And this is not to say that her reasons for believing it are just that good – no reasons are that good – but only that the mental state she finds herself in does not allow for the possibility of error. She has 'cognitive purchase on an objective fact.'[29] Allowing knowledge to be a mental state makes sense of the idea that we can be in a mental state that does not permit us to be mistaken. For an internalist, the mental state does not ensure the truth of what is believed because it is always possible for the belief, no matter how justified, to be false.[30] The external is irretrievably external.

Understanding knowledge in this way, as an external mental state, turns on its head the usual approach to epistemology. It is knowledge that justifies belief,[31] not the other way around. Knowledge is conceptually more fundamental than belief, is more basic than belief. What this means is that our direct relationship to the external world is not in question – beliefs about the world are not prior to knowledge about the world. We do not need to retreat to how things appear and attempt to work our way back to the world through reason alone. Instead we see our mental life as encompassing the world as it is, and this is what allows reason to function.

As McDowell has it, the possibility of appearances, the possibility of mental content depends on us having direct 'cognitive contact' with the world. Internalism, the idea that the mental (internal) realm is distinct from the external world, is not even intelligible:

> If moves in the space of reasons are not allowed to start from facts, riskily accepted as such on the basis of such direct modes of cognitive contact with them as perception and memory, then it becomes unintelligible how our picture can be a picture of a space whose positions are connected by relations that reason can exploit, such as that one of them is a reliable

[28] p.885.
[29] McDowell 1995 p.883.
[30] Excepting, perhaps, some *a priori* or self-evident truths.
[31] Williamson 2000, Chapter 9: He argues that one's evidence is what one knows. The concept knowledge is used to explain the concept of evidence, rather than, as is more usual, the concept of evidence being used to explain the concept of knowledge. (For a good discussion on making sense of this, see Hyman 2006.)

> ground for moving by inference to another. If the space of reasons as we find it is withdrawn from the objective world as it makes itself manifest to us, then it becomes unintelligible how it can contain appearances, content-involving as they must be, either. … *reality is prior, in the order of understanding, to appearance*, … it makes no sense to suppose that a space sufficiently interiorized to be insulated from specific manifest fact might nevertheless contain appearances.[32]

McDowell's point is that the idea of reason, as a mental process, does not make sense unless there is some kind of starting platform which is world entailing. The gap between justified belief (the internal) and truth (the external) will always ensure that the 'space of reasons'/the internal will have no purchase on the world. The ontological distinction presumed by internalism isn't just a problem for knowledge, it is a problem for the idea of there being a contentful mental space at all.

There is no scope to defend McDowell's point here.[33] In fact it has not even been my aim to defend externalism about knowledge – but only to make it more palatable. Satisfactorily defending externalism would be far beyond the scope of this thesis. However, espousing the view is essential to the success of my project, and thus I have attempted, in the above, to make it a plausible alternative to the status quo. Embracing externalism (if one is not already there) requires an overhaul of one's metaphysical commitments: it is a lot to ask. I have not spent time arguing for externalism, though there are many good arguments in its favour, but I have tried to make the metaphysical shift seem less dire. I have found that taking this view on knowledge makes so many seemingly inevitable epistemological worries seem not so inevitable any more. And one of these inevitabilities forced on one from an internalist perspective is the idea that the indirectness of testimony undermines or threatens the status of testimony-based beliefs, threatening their claim to knowledge.

[32] 1994 p.890 (emphasis added).

[33] Williamson defends a related point, arguing that we cannot make sense of rationality independently of knowledge, (2000, Chapter 8) and that 'contents ascribed in natural languages are broad' (p.53), i.e. mental content is world entailing. Williamson and McDowell both hold that understanding the mental as world-entailing is necessary if we want to understand the mental at all.

Section Two: The dissipatory approach

How does the externalist conception of knowledge help one to dissipate the problem of vulnerability facing testimony? According to Cassam's picture, dissipating an obstacle involves denying the existence of the obstacle, rather than attempting to overcome it. Overcoming an obstacle means denying that an accepted threat is incompatible with knowledge, dissipating it means denying that the obstacle exists. In this section I want to argue how embracing an externalist account of knowledge can allow one to dissipate the problem of testimony. I want to argue that the vulnerability claim - 'Merely understanding an assertion is insufficient grounds (inadequate justification) for coming to know what is asserted – need not be accepted.

This does *not* involve establishing how it can be that understanding an assertion *is* sufficient grounds (or adequate justification) for coming to know what is asserted: this would involve something like what the non-reductionists propose, that understanding an assertion is sufficient justification because we have, for example, an *a priori entitlement* to accept that what a speaker says is true. Arguing this is to accept that understanding *on its own* does not provide sufficient justification. It is the a priori connection between testimony and reliability (e.g. Coady), or between intelligibility and truth (e.g. Burge), that provides the justification/entitlement.

I intend to dissipate the obstacle by denying that there is an 'adequate justification' requirement on knowledge. The obstacle to knowledge, the vulnerability problem, is an obstacle *the existence of which* can be denied. If knowledge is not seen to be threatened by 'insufficient grounds', then the problem of vulnerability would not arise. It is not a matter of arguing that hearing an assertion does provide sufficient grounds, but of arguing that the problem of grounds need not arise in the first place. Dissipating the obstacle of vulnerability means denying that our testimony–based knowledge is threatened by inadequate justification. What will be denied is not that understanding an assertion provides insufficient grounds or justification for knowledge, but that knowledge requires sufficient grounds or justification.

The problem of vulnerability facing testimony seems inevitable because simply understanding an assertion that p is not sufficient to ground my knowledge that p.[34] I can understand an assertion, and yet it be false: it is possible for a speaker to make an assertion, and yet not believe it (she is lying); and it is possible for a speaker to make an assertion, believe it, and yet it still be false (she is mistaken). Hearing and understanding an assertion that p does not, by itself, get me to know that p because hearing and understanding an assertion does not provide me with adequate grounds for taking p to be true. This has prima facie plausibility, and yet, underlying this view is presumption that the possibility of false beliefs gained by testimony threatens the possibility of getting knowledge. The problem comes in, not because an assertion can sometimes be false, but because the possibility of an assertion being false *undermines the epistemic status of any belief we gain in this way*. It is not just to accept that one will *sometimes* not get knowledge from an assertion, but to accept that understanding assertions (on its own) *never* provides adequate grounds for knowledge as a result.

It is this move that the externalist can deny. The possibility of false testimony does not automatically result in a threat to all testimonial knowledge. The possibility of false testimony does not force us to maintain that understanding an assertion that p *underdetermines* knowledge that p. It is true that simply believing that p upon being told that p will *sometimes* result in coming to have false beliefs, but it is a further claim, and one that the externalist denies, to argue that this results in a threat to the status of all beliefs gained thereby. Why ought we to accept that the possibility of getting false beliefs from testimony automatically results in the claim that hearing and understanding an assertion is insufficient grounds for coming to know what is asserted? Why is it *inevitable* that the possibility of false testimony results in the claim that more than just hearing an assertion is required to get knowledge in this way? Why ought we to accept that the job of the epistemologist is to provide justification in order to establish that knowledge can be gained in this way?

[34] It is my intention to argue, in Chapters 4 and 5, that merely understanding an assertion *does* give one knowledge of what is asserted: I will argue that as long as the speaker makes a *correct* assertion, and the hearer understands it *correctly*, then the hearer will get knowledge of what is asserted – but this is not the same as arguing that understanding an assertion provides *sufficient grounds* for knowing what is asserted.

My proposed solution is to argue that the vulnerability problem, the insufficiency of grounds problem, is necessarily a problem only if one embraces the idea that knowledge requires justification (broadly construed) in order to have the status of knowledge. In this section I want to defend the idea that the externalist conception of knowledge can allow one to ignore this requirement. I will give reasons for thinking that an internalist conception of knowledge (the idea that knowledge is a state that reduces to internal and external components) will make it difficult for one to dissipate the problem of vulnerability, and that, in contrast, an externalist conception (the idea that knowledge is an irreducible mental state) can allow one to do just that.[35]

Espousing internalism about knowledge results in the possibility of error posing a threat to knowledge because the *possibility of error necessarily threatens the sufficiency of the grounds* one has for the belief, and thus the status of the belief qua knowledge, is undermined. If we reject internalism, then it is no longer inevitable that the indirectness of testimony will threaten the possibility of knowledge. Embracing the externalist conception of knowledge can be a way of avoiding the problem of vulnerability because externalism rejects the presumption that the possibility of false testimony necessarily results in an undermining of the status of testimony-based beliefs. It is the presumption that a belief's claim to knowledge is undermined by the possibility of false testimony, the idea that possibility of false testimony results in the threat to sufficiency of grounds, that makes the problem of vulnerability inevitable. The possibility of false assertion, or what is seen to be the 'flimsiness' of testimony as a source of belief, automatically threatens the status of beliefs *qua* knowledge only if said beliefs are dependent on the sufficiency of their grounds to establish them as knowledge. The idea that understanding an assertion is insufficient grounds for coming to know what is asserted arises because the possibility of the speaker lying or being mistaken undermines the possibility of understanding an assertion as being sufficient grounds for coming to know what is asserted.

The demand that is placed on justification by internalist accounts of knowledge is that it distinguish between accidentally true beliefs, and knowledge: just believing a speaker's assertion would not give one knowledge of what is asserted, because speakers lie and can be mistaken. The problem of vulnerability facing testimony-based knowledge arises

[35] Embracing externalism does not *necessarily result in* dissipation of the obstacle, but I want to argue that it can be used to dissipate it.

because there is a question as to whether beliefs formed on the basis of a speaker's testimony are formed in the *right way*. Of course it is possible to form true beliefs just by believing testimony, but in order for one to get knowledge from testimony, it must be shown that trusting the speaker is the epistemically right thing to do. There must be, in other words, a 'proper' epistemic link between hearing a speaker assert that p, and coming to know that p. The problem of vulnerability is the epistemic gap that exists between hearing a speaker assert that p, and coming to know that p. The hearer's getting knowledge from a speaker is dependent upon the process of testimony being justificatory, warrant providing, or epistemically correct.

This line of reasoning is only embarked upon because of the presumption that in order to determine a belief's status as knowledge, justification of some kind must be established. It is presumed non-negotiable that there has to be some kind of sufficiency of grounds in order for a belief to count as knowledge. But this presumption is not non-negotiable. In fact it is presumed false by externalists. Knowledge, according to externalism, is not a special kind of belief at all. Knowledge is a mental state that is distinguished from belief in its *factiveness*. Belief is conceptually subjective, it is a state of mind that does not necessarily say anything about the way the world is. Knowledge, on the other hand, is a mental state that incorporates the external world. It is a 'broad' mental state. It is a mental state that entails that the world is the way it is taken to be. Knowledge is not analysable into necessary and sufficient conditions, truth and belief making up two of those.

From the externalist perspective, it is possible that there is no epistemic gap to be crossed between hearing from a speaker that p and coming to know that p. As long as hearing that p from a speaker *is* a way of coming to know that p (the claim that I will defend in chapters 4 and 5) then, when one does get knowledge in this way, it is not possible that the speaker was lying or mistaken. Knowing that p (from testimony or any other source) entails that one cannot be wrong about p. One does not need to establish that one's belief is justified or that one is entitled to that belief in order to establish that belief as knowledge. The possibility of false testimony does not automatically result in a threat to the sufficiency of the grounds for knowledge, thus undermining the possibility of knowledge. If one gets to know that p from testimony, the risk of error (the risk of the speaker lying or being mistaken) no longer poses a threat to the status of that knowledge. The indirectness of testimony (the possibility of false testimony) does not result in threat to knowledge via a

weakening of the grounds or basis for accepting the word of another, because, when one does get knowledge from testimony, it cannot be that the speaker was lying or mistaken. Testimony-based knowledge is not undermined by the possibility false testimony, because having knowledge guarantees that nothing has gone wrong . The possibility of something going wrong means that sometimes one does not get knowledge when one thinks that one does, but this does not automatically result in a threat to the status of knowledge when nothing goes wrong.

The internalist view holds that knowledge is a state that is analysable: it is reducible to internal (or psychological) components (such as justified belief) and an external component (truth). The externalists, on the other hand, do not accept that knowledge is reducible to internal and external components, but hold that it is an irreducible mental state that encompasses the external world. The relevant point of distinction for me is that if one rejects the idea that knowledge has an external, non-mental component, then, when knowledge is had, it is no longer vulnerable to the threat of the possibility of error. This is because when one knows something, one cannot be in error. Knowing something *guarantees* that what one believes is true: there is no possibility of error. The internalist conception of knowledge, on the other hand, will be threatened by the possibility of error. This is because the possibility of error threatens *the sufficiency of the grounds* for the belief, and thereby threatens a belief's status as knowledge. The problem is that if assertions can give false beliefs, then merely understanding an assertion will never be enough (on its own) to provide knowledge.

Knowledge is a factive mental state:[36] the satisfactoriness of its epistemic position *qua* knowledge is not determined by the strength of the reasons or justification for the belief in question. Knowledge as a mental state gains its status of knowledge because it is *true*, and it being true is not something external and separate to the mental state itself: it partly *constitutes* the mental state of knowing. Knowledge is not a kind of belief, it is a mental state that entails an actual state of the world. Being in the mental state of knowing *guarantees* that what one knows is true. The mental state encompasses the truth, unlike for internalists, where the mental state is internal and distinct from the truth. For

[36] Williamson's phraseology.

internalists, the fact that truth is external to the mental state means that what establishes a belief as knowledge is the justification for the belief.

The externalist, it seems to me, can dissipate the obstacle of vulnerability by denying that the possibility of a speaker lying or being mistaken threatens or undermines the status of testimony-based knowledge. Knowledge is not belief with special epistemic status. Rejecting the internalist picture of knowledge allows one to dismiss the seemingly inevitable demand for justification. It is not justification that distinguishes true belief from knowledge, knowledge is not a special kind of belief with necessary and sufficient conditions. Knowledge is a mental state that encompasses the world being the way it is known to be. Knowledge gains its special standing because it is a *distinct kind* of mental state, different to mere belief: it is not a purely internal mental state, it is a mental state that entails some state of the external world. If one does get knowledge through testimony, then the speaker could not have been lying or mistaken.

But, the internalist will protest, how can I claim that the possibility of error poses no threat to knowledge? For surely, whether one is an externalist or an internalist about knowledge, if the speaker is lying or mistaken (her testimony is false) then one does not get knowledge of what she asserts. One cannot reason *from* the fact that one has knowledge *to* the fact that what one believes is true, for the very thing in question is how we can know that what we believe is true/is knowledge. Given that it is possible that a speaker could be lying or mistaken (and no one denies this) then a resulting belief's claim to knowledge is undermined. To be knowledge, it must be true, and thus, the internalist reasoning goes, in order to establish that we have knowledge, we must establish that the speaker is not (likely to be) lying or mistaken. But the externalist denies this move: yes, to be knowledge, it must be true, but no, this does not mean we have to establish that the belief is likely to be true in order to establish that it can be knowledge.

The externalist accepts that if a speaker is lying or mistaken, then knowledge cannot be had from her testimony. What is denied is that this fact threatens or undermines the possibility of knowledge in *every* instance of testimony. It is only when the speaker actually is lying or mistaken (when she does not know what she asserts) that the hearer does not get knowledge. When she is not lying or mistaken (if the speaker *does* know what she asserts), the hearer can and (if he accepts her testimony) *does* get knowledge. One does not need to

establish through justification that the speaker is not lying or mistaken, it only needs to be true that she is not lying or mistaken. The possibility of error threatens the possibility of *always being able to know that one knows*, it does not threaten the possibility that one knows when one does know. It is possible for one to sometimes think that one has knowledge, but not have it; but when one does have it, for example, when one knows that p, one cannot be wrong about p. The epistemic state of knowing is not undermined by the necessary indirectness of testimony. The possibility of a speaker lying or being mistaken does not result in a threat to the epistemic status of testimony-based knowledge.

The externalist rejects the idea that one must establish that the grounds for knowledge must be 'sufficient' in order for knowledge to have the status of knowledge. The externalist conception of knowledge does not put this kind of pressure on justification, because justification is not a necessary premise in an argument for the establishment of knowledge. One does not need to establish that the grounds for a belief are strong enough to come to conclusion that one knows. Instead, when one gains knowledge, the possibility of being wrong is excluded. One's grasp on the facts excludes the possibility of error, because grasping (knowing) the facts entails that they are as one takes them to be. What this means is that the indirectness of testimony (and the possibility of error) does not automatically undermine or bring into question the necessary conditions for the possibility of knowledge. When one does get knowledge from testimony, nothing has gone wrong.

This is not to say that *thinking* or *believing* that one knows means that nothing has gone wrong, only that *actually* knowing means that nothing has gone wrong. The internalist may baulk at this, claiming that if the two states (knowing, and thinking that one knows) are subjectively indistinguishable, on what grounds does one claim that nothing has gone wrong? But the fact that the two states are subjectively indistinguishable does not mean that the one's epistemic position (what one has epistemic access to) is the same in both. When one knows, nothing has gone wrong: one has a cognitive grasp on the way the world is. When one thinks that one knows but doesn't, not only has one failed to grasp the way the world is, one fails to grasp that one is not in the mental state one thinks one is in. The failure to know occurs within the mental realm. The idea that we need to be able to distinguish between knowing and thinking that one knows in order to know is not an idea that externalists (or many internalists, too, for that matter) accept.

Williamson's arguments on scepticism show that the possibility of error, the fact that something might go wrong, does not force one to retreat to the claim that it *only ever appears* to one that things are as they seem to be.[37] I contend that the same reasoning can be applied to testimony. The possibility of false testimony (such as when a speaker is lying or mistaken) does not mean that the hearer must retreat to the claim that all he has access to is that the speaker has asserted that p. Not getting knowledge when the speaker is lying or mistaken does not mean that one does not have knowledge when nothing has gone wrong. Instead it means that one's failure to know when something has gone wrong is a failure on two counts: it is a failure to know, and a failure to know what mental state one is in: 'Part of the badness of the bad case [when one does not know that one does not know] is that one cannot know how bad one's case is.'[38] When one is in the 'good case' (nothing has gone wrong) one succeeds in knowing, and one also knows that nothing has gone wrong.

Something like this is what McDowell means when he defends the claim that we ought to allow luck to 'enter the space of reasons.' The unluckiness involved when something goes wrong, as, for example, when a speaker is lying or mistaken, is something that occurs within the space of reasons. Internalists hold that luck is a part of the external world, that when we have knowledge reason or rationality does its best, and whether or not what is believed is true is down to the world cooperating. On McDowell's view, however, the 'cooperation of the world' is something that happens within the mental realm: 'the particular facts that the world does us the favour of vouchsafing to us, in the various relevant modes of cognition, actually shape the space of reasons as we find it. … Of course we are fallible in our judgements as to the shape of the space of reasons as we find it, or – which comes to the same thing – as to the shape of the world as we find it. That is to say we are vulnerable to the world's playing us false; and when the world does not play us false we are indebted to it.'[39]

---

[37] Chapter 8.

[38] 2000 p.165, and his 'Scepticism and Evidence' (2000).

[39] 1995 p.887. When he says we are indebted to the world, he does not mean that, as the internalists have it, whether or not we have knowledge is up to the world. What he means is that when we *judge* that we have knowledge, we are indebted to the world: 'we are fallible in our judgements as to the *shape of the space of reasons* as we find it' (emphasis added).

The picture being painted here is of knowledge being fallible, yet not in the usually accepted sense that justification cannot guarantee the truth of our belief, but rather that we are fallible in our judgements about whether we have knowledge or not. This does not mean that we can *never know* that we know something, only that knowing something does not require that we are able to know that we know. It is possible to think one has knowledge, and yet not have it. Understanding this distinction is essential to understanding how it is that externalism can allow one to avoid the problem of vulnerability posed by the indirectness of testimony. The indirectness of testimony poses a threat to knowledge *only if* justification is meant to establish a belief as true (as knowledge). If knowledge is something that does not depend on justification or reason for its status, but instead is a mental state that is partly constituted by the fact itself, then the possibility of error does not undermine it in the same way. It is possible to think that a speaker's assertion is true when it is not, and to think that one has knowledge when one does not, but this reveals only that we can sometimes be wrong about what we think we know, and *not* (as the internalists presume) that the reason or justification we have for thinking she is speaking the truth is insufficient to establish that we know it.

The possibility of error does not undermine the justification we have for our beliefs, thereby threatening the possibility of knowledge. What goes wrong when we think we have knowledge but do not (when we accept false testimony) is an error within the 'space of reasons.' There is no epistemic gap created thereby between how things appear, and how things are (between the internal mental state and the state of the world). When we have knowledge, then things are necessarily how we take them to be, we cannot be wrong. How things appear is not separable from how things are. The possibility of error enters at a different level, at the level of our ability to always be able to tell what mental state we are in. Sometimes one can be wrong about what mental state one takes oneself to be in. When everything goes right, and I know that p, then it is not possible for me to be wrong about p. I can know that p, and I can know therefore that the speaker was not lying or mistaken. When everything goes right, and I get knowledge that p from a speaker who tells me that p, I cannot be wrong about p. It is not possible that the speaker was lying or mistaken.

Internalists, like externalists, hold that knowledge entails that what is believed is true (knowledge is not just the internal mental components, what is believed must also be true). But for internalists this is just to say that, *by definition*, when a hearer gets knowledge, the

belief must be true; it does not explain how knowledge can be had. The problem of vulnerability automatically arises for internalists because the possibility of the speaker lying or being mistaken threatens a belief's claim to knowledge. The indirectness of testimony (the possibility that a speaker could be lying or mistaken) undermines the epistemic status of the belief in question, and how he gets knowledge in the face of this is a problem that internalists have to deal with. For the externalist, the grounds that one has for dismissing the possibility of something having gone wrong is that the hearer knows that p. The hearer is in a mental state that, unlike mere justified belief, *guarantees* that what he takes to be the case, is in fact the case.

Externalists hold that the truth of what one believes when one knows that p is not something that is determined by a fact external to the knower. The truth of what one knows is part of the mental state of knowing, which means, when something goes wrong and one thinks that one knows, and yet does not, one believes oneself to be in a mental state that one is not in. When one fails to know (but still thinks that one knows), it is not that the external world has failed to coincide with one's beliefs, but rather that one fails to recognise what mental state one is in. The truth or falsity of the belief is not something that occurs only in the world, it occurs within the mental realm. When one accepts a speaker's testimony, one can come to have knowledge (if everything goes well), but one can also fail to get knowledge (if the speaker, unbeknownst to the hearer, is lying or mistaken.) For the externalist, in the latter case, one fails to recognise that one is not in the mental state of knowing (he thinks that he knows, but does not), but in the former case one can be in the mental state of knowing, without establishing that nothing has gone wrong – without having justification that overcomes the threat of false testimony.

The fact that a speaker could be lying or mistaken is a risk that other sources of knowledge do not carry, and yet this does not mean, as internalists hold, that we ought, therefore, to have a justification (or a story about reliability or entitlement) for taking a speaker at her word. The possibility of the speaker giving false testimony does not undermine the status of the knowledge gained thereby. But the claim that sufficient justification or grounds is not a requirement on knowledge is not the same as claiming that gaining knowledge through testimony has nothing to do with being rational and responsive to reasons. It is not to say that we should just believe everything we hear. We must be what McDowell calls

'doxastically responsible'[40] – we must be sensitive to the possibility of things going wrong. But being doxastically responsible and rationally aware of how things can go wrong is not the same as making a justification requirement on knowledge through testimony.

An irrationally formed belief, even if true, ought not to count as knowledge. However, although the externalist can accept that knowledge can only be knowledge if it is formed rationally, it does not automatically follow that in order to establish that one has knowledge, one must establish that the belief in question is rational, or based on a reliable source. I can know the capital of Sudan without having any idea about how I know it. Much, if not most, of my knowledge of historical, geographical, economical facts, etc. is knowledge that I do not question, and yet could not say if the source was reliable. It might be safe to presume that I must have heard it from someone trustworthy, or read it somewhere reliable, but this would be a presumption based only on the fact that I count myself as knowing these things. I do not need to establish that the source of my knowledge was reliable, or that the belief is rational in order to establish that I know that the capital of Sudan is Khartoum. And in the case of testimony, one does not need to establish why we are justified or a priori entitled to accept it in order to establish that we can get knowledge in this way.

It is not controversial to say that an irrationally formed belief is not knowledge, but rather than placing a demand on justification such that it needs to establish a belief as knowledge, the relationship works the other way: the fact that one knows something explains why one's reasons for believing it are good, or ensures that the belief was rationally formed. Knowing happens in the 'space of reasons' (the mental realm): being rationally sensitive or 'doxastically responsible' forms the necessary background upon which knowledge can occur, but it is not the job of reasons for belief (entitlement etc.) to distinguish mere true belief from knowledge. Knowledge is not reducible to a kind of belief that is true and is justified. In the case of testimony, we must be sensitive to the possibility of a speaker lying or being mistaken: if there are reasons to think something has gone wrong, then we cannot get knowledge. But this is the rational background upon which we can gain knowledge from a speaker, it is *not the condition by which we judge whether or not we have knowledge*.

[40] McDowell 1994 p.210.

Knowledge is a *mental* state, it is a state that occurs within the 'space of reasons.' What this means is that its being subject to reason is constitutive of it counting as knowledge. As McDowell puts it: 'one cannot count as occupying an epistemic standing with respect to a proposition [cannot count as knowing it] unless, in taking things to be so, one is responsive to what is probable given one's informational position.' [41] Our responsiveness to reasons is what qualifies knowledge as a state that occurs within the 'space of reasons'. Doxastic responsibility, responsiveness to reason or rationality are not a condition on knowledge in the way that the internalist view on knowledge proposes. It is not the case that knowledge can be analysed in terms of rationality or doxastic responsibility. The fact that knowledge occurs on a background of rationality or responsiveness to reason is what allows us to be knowers. Responsiveness to reasons or probability forms the 'necessary background for talking of positions in the space of reasons at all.'[42]

What this means is that justification, reliability or stories about epistemic entitlement cannot be used to distinguish between knowledge and mere true belief. Establishing reliability is not a necessary or sufficient condition for establishing knowledge.[43] What makes knowledge distinct from mere true belief on the externalist account is that a mere true belief is a composite state. The states are metaphysically distinct. As Williamson has it: 'Even if believing truly is a mental state in some liberal sense of the latter term, there is also a more restrictive but still reasonable sense in which believing truly is not a mental state but the combination of a mental state [believe] with a non-mental condition [truly]. The present claim is that knowing is a mental state in *every* reasonable sense of that term.'[44] The externalist resists the move that knowledge must be analysable into more basic concepts: knowledge is more basic than belief. An accidentally true belief is not *not knowledge* because it lacks justification or reliability, or is accidentally true, it is just not knowledge – it is a metaphysically distinct state.[45]

---

[41] 1994 p.209.
[42] p.210.
[43] Turri (2013) argues that knowledge can even be achieved from an *unreliable* source. This may be true, but my claim in the above is not that reliability is not necessary for knowledge, but that *establishing* reliability is not necessary for establishing knowledge. If reliability *is* necessary we can use the fact that we have knowledge to determine that the source was reliable.
[44] 2000 p.28.
[45] Whether a belief is accidentally true or 'non-accidentally' true (e.g. based on justification/reliable source etc.), makes no difference if the state it is meant to establish is knowledge. Neither accidentally true nor 'non-

Rationality or responsiveness to reasons does not say anything about knowledge itself, false beliefs can be rational too. Rationality or responsiveness to reason is not a condition which can be used to distinguish knowledge from other beliefs, it is not a necessary condition for knowledge in the way that the internalists use it, it just forms a necessary background if knowledge is to count as being a state that occurs within the space of reasons. So although it is true that knowledge requires responsiveness to reasons, this does not mean that the indirectness of testimony, and the possibility of a speaker lying or being mistaken, threatens or undermines the possibility of testimony-based knowledge. If a hearer has reason to think that a speaker may be lying or mistaken, then he cannot get knowledge from her – it would be irrational to accept what she says as true. But this does not mean that he has to have reason (or justification more broadly) to accept that she is *not* lying or mistaken in order to get knowledge from her. Nor does the epistemologist need to establish that testimony is a reliable source of knowledge (or that we are entitled to rely on testimony) in order to establish that we can get knowledge in this way. There is *no need to provide a justification* in order to establish that knowledge is possible.

When we know something, what we know is *directly manifest* to us. The fact itself forms part of the mental state of knowing. There is no gap between how things appear to us, and how they are in reality, a gap that must be crossed with some justificatory story. Knowing something entails that one is in a mental state such that what is believed *cannot be* false. The problem of vulnerability posed by the indirectness of testimony presupposes this gap between how things appear (a speaker asserting that p) and how things are (the fact that p). This gap threatens the possibility of testimony-based knowledge because hearing a speaker assert that p does not seem (at first sight) to be a good enough reason to ground knowledge that p. This gap is meant to be overcome via provision of a justification.

The problem can be dissipated, however, if we deny that there is an epistemic gap to be crossed. Knowing that p based on hearing a speaker assert that p *forbids the possibility that p is false*. The fact that p is made manifest to the hearer upon hearing it asserted that p. There is no epistemic gap to be crossed. It is not the job of the epistemologist to justify or give epistemic support for the beliefs that we get through testimony: so doing (as I

---

accidentally' true beliefs are knowledge. But knowledge, as the more basic concept, can make sense of what we mean by justified or 'non-accidentally' true. (See Williamson's arguments for the claim that knowledge *is* evidence, 2000, Chapter 9.)

suggested in the previous chapter) will not provide a completely satisfactory answer to the question of how testimony-based knowledge is possible. If there is an epistemic gap between hearing that p and coming to know that p, this gap cannot be completely overcome. The best way to deal with the problem is to deny that it exists.

One can deny that the possibility of false testimony results in the threat to the sufficiency of the grounds one has for coming to know what is asserted by denying that knowing the proposition asserted requires 'sufficient grounds' for taking it to be true. Instead, one can come to know what is asserted just by hearing and understanding an assertion. This is not to say that understanding an assertion provides *sufficient grounds* or *adequate justification* for coming to know what is asserted, but that understanding an assertion is, more basically, *a means* to coming to know what is asserted. One can come to know that p upon hearing and understanding an assertion that p, because hearing and understanding assertions is a way of coming to know what is asserted.

If, however, there is to be *no epistemic gap* between hearing (and understanding) an assertion that p, and coming to know that p, it must be the case that simply understanding an assertion that p *is a means to* coming to know that p. If no *justification* or *grounds* are necessary (or sufficient) to explain how knowledge can be had through merely understanding an assertion, then merely understanding an assertion must be a sufficient *explanation* for the possibility of gaining knowledge in this way. In order to understand how we can get knowledge in this way, we need to know how it is that understanding an assertion that p can get one to know that p, without recourse to any justificatory story.

It needs to be the case that when a hearer comes to know that p upon hearing (and understanding) an assertion that p, the fact that p is *made manifest* to him. He does not need an epistemic warrant or entitlement to justify his conclusion that p is true, based on hearing the assertion that p. He does not need to work from the premise that the speaker has asserted that p to the conclusion that p is true. When he properly understands the assertion that p, he apprehends, thereby, that p is true. The truth of p is communicated to him directly through the assertion employed by the speaker. This is the picture that I aim to defend in the next two chapters.

Conclusion

This chapter constitutes the second part of the obstacle removal stage of Cassam's three tiered approach to how-possible questions. In the previous chapter I argued that there was reason to think that an overcoming approach to the obstacle, the approach that most epistemologists in the field take, might not be completely satisfactory. The fact that justified belief is distinct from knowledge means that establishing justified belief is not tantamount to, or equivalent to, establishing knowledge. The gap between belief and knowledge is one that cannot be crossed by justification alone. The difficulty is that seeing the problem as one of vulnerability (that is, that the indirectness of testimony poses a threat to the possibility of knowledge) undermines one's ability to overcome the problem.

The purpose of that chapter (Chapter 2) was to motivate for taking a different, namely dissipatory, approach to the problem. If overcoming the problem of vulnerability is not (or might not be) possible, then there is good reason to attempt to dissipate it instead. In this chapter I argued that the commonly accepted internalist conception of knowledge made such a response difficult. If knowledge is seen to be a state that is comprised of internal and external components, then seeing the problem of testimony as a problem of vulnerability is inevitable. This is because knowledge, under this view, depends upon its reasons or justification to establish it as knowledge, and the possibility of a speaker lying or being mistaken threatens the reasons or justification we have for taking a speaker to be telling the truth. Instead of providing justifications, I argued, we should reject this view of knowledge.

The externalist conception of knowledge, as distinct from the internalist, holds that the mental realm, or the mind, or the 'space of reasons', is a space that encompasses, or is partly constituted by, the external world. What this conception of knowledge allows is that knowledge qua knowledge does not depend on justification for its status *qua* knowledge. The fact that it occurs within the 'space of reasons' ensures that rationality or responsiveness to reason form a necessary backdrop to knowledge, but what makes it knowledge is that it entails a state of the world. The possibility of error that the indirectness of testimony implies is not a threat to knowledge, because such a possibility does not and cannot require that we retreat to the claim that all we have access to is what appears to be the case (that the speaker has asserted that p). Hearing a speaker assert that p gives us

access to the fact that p. The indirectness of testimony does not undermine the epistemic status of testimony-based knowledge.

Perhaps at the forefront of the reader's mind is the question of *how* is it that a speaker's assertion that p can make it directly manifest to the hearer that p. If knowledge requires a factiveness such that the truth of what is known forms a part of the knower's mental landscape, how can such a state be communicable? What is it about making an assertion, and hearing an assertion that results in there being *no epistemic gap* between hearing that p and coming to know that p? In other words, how is testimony-based knowledge possible? Cassam's structured approach to how possible questions has a third level, the level of enabling conditions. This is the level at which a richer, more satisfactory response to the question of how is testimony-based knowledge possible can be put forward. This is the level of *explanation*.

The remainder of the thesis is devoted to this task. In Chapters 4 and 5 I defend the claim that properly understanding an assertion entails that the hearer comes to know what is asserted. In Chapter 4 I argue that communication as an institution is one that allows a speaker to communicate to a hearer her actual mental state. There is no epistemic gap between properly understanding a communicative utterance and coming to recognise the speaker's actual mental state. In Chapter 5 I argue that the mental state communicated in the instance of testimony is knowledge. Knowledge is a factive mental state, and the factiveness (the truth of what is asserted) is available to the hearer upon understanding the assertion. In this way I hope to explain how it is that knowledge can be transmitted from a speaker to the hearer. If it is possible for mere communication to be the means of transmitting knowledge, then the possibility of communication explains the possibility of testimony-based knowledge.

Chapter Four: Communication and factive understanding

In the previous chapter I argued that the problem of vulnerability facing testimony ought to be dissipated rather than overcome. The indirectness of testimony does raise questions of how testimony-based knowledge is possible, but not because indirectness threatens or undermines a belief's claim to knowledge. We do not, I argued, need to answer the question of how testimony-based knowledge can have the epistemic status of knowledge, because knowledge is not a mental state that depends upon its justification to establish it as knowledge. Knowledge, qua knowledge, is not undermined by the threat of the speaker lying or being mistaken in the way that it is often presumed. However, this is not to say that knowledge can just be had from testimony, and there is no longer a question of how it is possible. There is still a question of how testimony-based knowledge is possible, but it is a substantial question about the process of testimony and how it works as a way of transmitting knowledge.

The question of how testimony-based knowledge is possible, I believe, should not focus on an attempt to overcome the problem of vulnerability with a story about how we are entitled to or justified in our testimony-based beliefs. Instead, the focus should be on how exactly testimony works, and explaining how one person can share or transmit her knowledge of something to a hearer via a communicative utterance. This type of explanation constitutes the third level of Cassam's three tiered approach to how-possible questions. It is the level at which positive or substantive explanation can be put forward. The aim is no longer to defend the possibility of testimony-based knowledge from the threat of vulnerability, the aim is no longer to show how it is that testimony is *not im*possible. The aim is to establish what Cassam calls the 'enabling conditions.' The enabling conditions are the conditions which need to be in place for knowledge transmission through testimony to occur,[1] they are the conditions that *explain* the possibility of testimony-based knowledge.

[1] The institution of communication is not the *only* necessary condition for the transmission of knowledge, (see Chapter 1, n.4 and n.20), if communication is to work, other conditions must be in place. For example, there must be a speaker, a hearer, a context, a common language etc. But if communication does occur, then these other conditions will be in place too. It is the institution of communication, and not the other necessary conditions, that *explains how* knowledge is transmitted through testimony.

When a person gains knowledge through testimony, the process through which this occurs is one of communication. The speaker has knowledge, and shares this knowledge with the hearer via communication. But communication does not always involve the sharing of knowledge; communication is also about letting a hearer know of one's *internal* mental state. Communication between humans (usually, but not always, linguistic) is what happens when information (broadly construed) is passed intentionally from a speaker (or speakers) to a hearer (or hearers). Typically, the speaker uses an utterance to communicate with the hearer. The speaker can communicate to the hearer an emotion, an attitude, a belief, an opinion, an intention, a desire, a thought, a wish, a feeling; speakers can ask for something, make commands, requests, promises. And speakers can also share their *knowledge* of states of the world with hearers.

When a speaker successfully communicates with a hearer, she communicates her mental state to him: he recognises her as having, or being in, a certain mental state – the one that she actually is in. Properly understanding a speaker's utterance gives one direct access to the speaker's actual mental state – the hearer knows what mental state the speaker is in. When knowledge is transmitted (that is, in the case of testimony) more than this happens: the hearer not only recognises the speaker as having a certain mental state, but he comes to have it too. Knowledge is contagious in a way that other mental states are not, and this requires further explanation. My belief is that if we understand knowledge to be an irreducible external mental state, in the way I have been promoting, we can understand the factivity of the mental state of knowing as playing a key role in this explanation. Defending this will be the aim of the next chapter, however. The aim of this chapter is to defend the claim that successful communication of mental states entails that the hearer recognises the speaker's actual mental state.

In order to explain how this happens, I am first going to present my view on how communication of internal mental states succeeds. This is the aim of this chapter. (The account will avoid, at this point, talk of assertions and communicating states of affairs as opposed to internal mental states.) I hope to defend the idea that communication, properly understood, has a necessary connection with factivity: the speaker's *actual* mental state (whatever it is that she is trying to communicate) plays a central role in explaining the process of communication. In Section One I will present my account of communication, an account that proposes that the possibility of communication depends upon two constitutive

rules: one rule that guides the behaviour of speakers (how to make utterances), and one rule that guides the behaviour of hearers (how to understand utterances). In Section Two I will attempt to illuminate how my 'factivity entailing' account is better able to explain the phenomenon of communication than the more popular 'thinner' conceptions of communication.

One thing that should be noted here is that the account that I defend of communication is only sketched out here. There is no space in this thesis to develop it properly, and to consider the inevitable complications that will arise out of any discussion on something so rich and complex as human communication.[2] I hope that I can do enough to defend the main tenets of the view, and give sufficient reason to think that they illuminate something important about the process of communication. However, how this account affects (and is affected by) important issues such as linguistic meaning, metaphor and sarcasm, and all of the subtleties involved in human interaction, for example, will not be considered. I just hope that concerns about these issues will not impact too heavily on the plausibility of the core claims that I make below.

Section One: Two rules of communication

The way that communication is made sense of and expanded upon varies greatly across and within disciplines and the literature on this subject is dense and prodigious. Situating my account within this debate or collection of debates would be difficult, and would obscure, rather than illuminate, my central claims. For this reason, I am first going to propose my account of communication without using any recognised framework or particular terminology in order to avoid coming into direct conflict with any extant theory on communication. However, that said, much of what I say may sound familiar as my account has grown out of my acquaintance with speech act theory (especially John Searle), communication theory (Bach and Harnish's 'Linguistic Communication'), and Gricean accounts of meaning, etc.[3] In Section Two I try to differentiate my account from these

[2] The proliferation of footnotes will be testament to this fact.

[3] Also very influential for me was a power point presentation of Peter Pagin's called 'When does communication succeed?' (Stockholm 2012); his paper 'What is Communicative Success?' (2008); Zeno Vendler's 'Telling the Facts' (1979); and H.H Clark's *Using Language* (1996). John McDowell's paper 'Meaning, Communication and Knowledge' 1980 played a large role in how I think about communication. In this paper

'thinner' accounts of communication, and defend it as explanatorily superior. My aim in *this* section is to connect with our pre-theoretic intuitions about communication, and to make sense of how communication (of, for now at least, internal mental states) works.

I will say, however, that in the extant theories, there is a general tendency to focus on the *speaker's* action, and to view the hearer's part in the process as a passive one. All that the hearer needs to do is 'understand' what is being said or communicated, and in so doing, communication is achieved. Just what this understanding entails is mostly ignored. But communicative success depends on the hearer 'getting the message', or being successfully 'communicated with'. And this is a very specific kind of understanding – it involves more than just understanding what the utterance means, it involves recognising something in the speaker. In this way my account aims to be more balanced than the majority of work done on speech acts and communication, it focuses on the role of the hearer as much as on the role of the speaker in the achievement of communicative success.

Linguistic communication is a two-way street. Communicating is not something that a person can do on her own, it requires a speaker *and* a hearer.[4] There must be, by definition, something (information in a broad sense) that is communicated, and to be communicated, it must pass from one person to another person. The message has to be 'sent' and 'received', as it were. H.H. Clark draws an analogy between using language and playing a duet, or dancing a waltz.[5] It is a 'joint activity',[6] an activity that can only be understood as something that people do *together*. One participant's actions cannot be explained or understood without reference to the other participant's actions, they are *co-dependent*.

My account is an attempt to spell out this 'joint activity', to explain how it works. It is my contention that the hearer is not a passive recipient of information (as is often presumed), but that the ability to understand depends upon there being a constitutive rule guiding one's interpretation of the speaker's utterance. It is not controversial to claim that

---

McDowell argues against Strawson's conception of communicative intention, arguing that 'the appropriate mutual awareness is actually what is aimed at by the speaker's primary communicative intention, so that securing the mutual awareness is not, as in Strawson's picture, a fallible means to communicative success, but rather constitutes it.' (p.130) McDowell does not talk of constitutive rules, and likens intentional communication to more instinctual or animal communication, but I believe that the conclusion I am ultimately aiming for is not dissimilar to McDowell's.

[4] 'Speaker' and 'hearer' are used as very general terms to describe those 'making' the communication, and those receiving it. (I will refer to a speaker as 'she' and hearer as 'he' throughout the thesis.)

[5] 1996 p.3.

[6] *Ibid.*

making utterances is rule guided,[7] but it is a little more so to claim that understanding utterances is also rule guided.[8] The ability to make utterances and to understand them places constraints on the speaker as well as on the hearer. It is my contention that these constraints, these constitutive rules, are what give rise to the possibility of communication.

The idea that communication is an institution sustained by the existence of constitutive rules is quite a familiar one. John Searle is one of the first to defend the idea that linguistic communication depended on constitutive rules: 'the semantic structure of a language may be regarded as a conventional realisation of a series of sets of underlying constitutive rules, and that speech acts are acts characteristically performed by making utterances in accordance with these sets of constitutive rules.'[9] I want to continue this legacy, but not in the same way as Searle.[10] I want to defend the claim that communication as an institution is sustained by the existence of (only) two distinct rules: a rule that governs the communicative behaviour of the speaker (how to make a communicative utterance), and a rule that governs the behaviour of the hearer (how to understand a communicative utterance).

But firstly, more must be said on what I mean by constitutive rules. Constitutive rules are to be distinguished from regulative rules.[11] Regulative rules are rules that govern a certain pre-existing kind of behaviour, such as good table manners, or school rules. Constitutive rules are rules that create or maintain an otherwise non-existent form of behaviour, and here the most common example cited is the rules of a game. The game itself, and the behaviour of the players or participants, is only explicable in terms of the rules that govern them. The action of placing your opponent in check-mate in a game of chess is dependent

[7]Using language, and communication in general, are usually taken to be rule guided activities, though how this is cashed out varies widely.

[8] I have not read anyone who makes this claim, though perhaps Angus Ross comes closest when he claims 'It follows that to knowingly and openly perform an action that is permissible only if a certain condition obtains is to place witnesses under a *prima facie obligation* to assume that the condition in question does indeed obtain' (1986 p.78 emphasis added). Though for Ross this obligation comes from our responsibility as 'full, adult' members of a community who have 'respect for others as fellow judges', not from a fundamental rule that sustains the very practice of communication.

[9] 1970 p.37.

[10] Searle defended the idea that each type of speech act had its own constitutive rules, whereas I defend the idea that there are two constitutive rules more basic than the different kinds of speech acts, rules that explain all communication.

[11] There is a good discussion on constitutive rules in Williamson (2000) pp.239-241, see also Cherry (1973), and Hindricks (2009).

on the existence of the rules of the game. Such an action is not possible outside of the existence of the rules for a game of chess.

Constitutive rules succeed in creating or maintaining a kind of behaviour because the players, or participants of that behaviour , *know the rules*.[12] Players of a game of chess count as players of chess because they know the rules of the game. Someone who moved pieces randomly around the board, even if these movements happened to be in accordance with the rules, would not count as playing chess. She has to know that the rules of chess govern her actions (even if she chooses to break those rules). Knowing these rules is synonymous with knowing how to play: she *understands what she is doing* when she moves her rook, and she understands what her opponent is doing when he moves his bishop. The fact that the rules govern the participants' behaviour is what creates the game, makes it possible, essentially *is* the game.

Constitutive rules are necessary for the existence of the institution in question. They are necessarily *in force* for participants in the institution. This is not to say that the participants must always obey the rules, constitutive rules can be broken or flouted.[13] But in order to even count as participants in the type of behaviour in question, their actions must be *governed by* these rules. The rules stand, whether they are broken or obeyed. The participants must be aware of the rules in such a way as to make their behaviour explicable. They need not be aware of them in the sense that they could write them down, or even be aware of them in the sense that they know that they exist. But their very participation in the institution depends on them being aware (in whatever sense necessary) that they are subject to those rules.

The reason that constitutive rules are called rules (as opposed to convention) are because they can be broken, behaviour can conform to or flout the rules.[14] Rules govern behaviour, they do not simply describe it. However, rules imply normativity, and it is important to be clear about the kind of normativity involved in constitutive rules. Normativity is a concept that often implies sanction, blame, and culpability, and yet this is

[12] 'Knowing the rules' is meant in a thin, non-explicit sense, that should hopefully become clearer as this chapter progresses.
[13] This is something that Searle's account seems to miss, he defines constitutive rules in such a way as to make it almost conceptually impossible to break them.
[14] Although Searle's account of constitutive rules does not seem to entail this – his essential rules for speech acts are 'counts as' rules: '*X counts as Y in context C.*' (p.36) and thus are not normative.

not the kind of normativity entailed in constitutive rules. (It is the kind of normativity entailed in regulative rules.) Constitutive rules are normative in that they *define* what is correct behaviour, and what is incorrect behaviour, but this does not necessarily result in sanction or blame. Constitutive rules are not so much *have-to* rules, as they are *how-to* rules. They do not come with *obligations*, so much as guide-lines that define the behaviour in question. There need be no moral aspect whatsoever to constitutive rules. The point of them is to define or create a kind of behaviour, and knowing the rules is synonymous with *knowing how to* participate in the institution in question. In games one is often sanctioned for cheating (breaking the rules), and yet this is not a *necessary* feature of playing games. Perhaps sanction plays an instrumental role in maintaining fairness, for example.

Constitutive rules do not need to feature in our consciousness, we do not need to be aware of them *qua* rules. One could be oblivious to the idea that one's behaviour is governed in this way.[15] Constitutive rules do not need to regulate our behaviour in a moral sense, the correctness or incorrectness of certain actions may be of no real consequence. We may break a constitutive rule over and again without people caring or even noticing. However, this does not interfere with the centrality of their role. Constitutive rules *create* a kind of behaviour, the existence of these rules is *one and the same thing* as the existence of the game or institution, and knowing how to play the game or participate in the institution is the same thing as knowing the constitutive rules that govern one's behaviour.[16]

To return to the rules for communication: it is my contention that the possibility of communication as an institution depends on the existence of two central rules, one rule governing the speaker, and one rule governing the hearer. Participants in the institution of

[15] Another example of this kind of rule governed behaviour could be how to pronounce words when reading aloud. Knowing *how* to pronounce the word 'red' is the same thing as knowing that one *ought to* make certain sounds upon seeing the letters R, E and D. One need not be aware that one's behaviour is rule governed, but the fact that one knows how to read aloud entails that one recognises that there is a *correct* way to pronounce words.

[16] How to make sense of the idea of rules governing behaviour in this way is a philosophical issue in itself. Is it necessary to know the rules in any kind of explicit way in order for one's behaviour to be explicable only in terms of rules? I hope not: 'We saw that a rule, properly speaking, isn't a rule unless it lives in behaviour, rule-regulated behaviour, even rule-violating behaviour. … In attempting to grasp rules as rules from without, we are trying to have our cake and eat it. … A rule is lived, not described.' (Sellars in Brandom 1994, p.25.) Without getting into the metaphysics of rule following, I hope that I can coherently hold that communicators need not be *aware* of the rules in any explicit sense, and yet their communicative behaviour be necessarily rule governed.

communication know how to communicate: the speaker's rule explains how speakers know how to make utterances, and the hearer's rule explains how hearers know how to understand them.

> **Speaker's rule**: The speaker's utterance[17] must be an explicit[18] attempt to get the hearer to recognise her (as having an) intention/desire/belief/attitude (or relevant mental state),[19] the i/b/d/a that she actually does have;
>
> **Hearer's rule**: The hearer must recognise, upon hearing the speaker's utterance, that it is an explicit attempt to get him to recognise her (as having an) intention/desire/belief/ attitude, the i/d/b/a that she *actually does have*. (This rule entails factive understanding: the state he takes her to be in must be the state that she is trying to get him to recognise her as being in AND she must actually be in it.)

The speaker's rule governs utterances by identifying utterances that are attempts to get the hearer to recognise[20] the actual mental state of the speaker as *correct*, and identifying utterances that do not attempt to get the hearer to recognise the speaker's actual mental state as *incorrect*. The hearer's rule governs interpretation by identifying interpretations that consist in recognition of the speaker's utterance as an attempt to get him to recognise her actual mental state as correct, and interpretations that do not as incorrect. Only if *both* rules are conformed to is communication correct, or technically successful. The rules could

---

[17] The utterance is the communicative act that the speaker employs to get the hearer to recognise her intention/belief/desire/attitude, and does not have to be linguistic (e.g. nods, shrugs, non-conventional signs).

[18] By 'explicit' I mean to capture what speech act theorists call the 'reflexive' nature of communicative intentions: the speaker's must want to hearer to recognise their communicative intentions – what they are attempting to do with the utterance. Searle, in 'What is a Speech Act' (1965), highlights the importance of this explicitness: 'I might on the one hand attempt to get you to believe that I am French by speaking French all the time, dressing in the French manner, showing wild enthusiasm for de Gaulle, and cultivating French acquaintances. But I might on the other hand attempt to get you to believe that I am French by simply telling you that I am French.' (p.7) The speaker, with her utterance, must be trying to get the hearer to recognise her mental state, *via his recognition of this being what she is trying to do*.

[19] This list is meant to capture the kinds of things that people communicate to each other. People want to know things, desire that a person do a certain thing, have certain attitudes, feelings or opinions that they wish to share with or reveal to the hearer. The kinds of things that are communicated are as complex and varied as the human experience.

[20] Although one might think 'reflect' or 'express' might work just as well here as 'is an attempt to get the hearer to recognise her as having', this is not so. 'The speaker's utterance reflects her actual mental state' sounds better than 'The speaker's utterance is an attempt to get the hearer to recognise her actual mental state.' But the utterance, for me, is not a representation of the speaker's mental state, but *constitutes* the speaker's communicative intentions – what she is attempting to do, no more and no less. What the hearer recognises, when he understands an utterance, is not what the utterance *represents*: what he recognises is the speaker's attempt to get him to recognise her i/b/d/a.

be called 'factivity entailing': a successful instance of communication entails that what is communicated is the speaker's *actual* mental state. *There must be a match or harmony between the speaker's actual mental state, the state that she is trying to get the hearer to recognise her as having, and the state that he thereby recognises her as having*.

These rules define technically correct communicative behaviour, and instances of communication that do not conform to the rules are still instances of communication, they are just necessarily flawed. It is possible for communication to occur when the speaker has hidden intentions, but it is necessarily incorrect communication, the speaker is breaking the constitutive rule. And it is possible for communication to occur when the hearer does not recognise the speaker as trying to get him to recognise her as having the i/d/b/a that she actually does have, (he may reject her utterance as an attempt to get him to recognise her actual mental state, or he may think he is recognising her actual mental state, but not be), but this too is necessarily incorrect, the hearer is breaking his constitutive rule. However, although the rules can be broken, they always *stand* in the sense that communicators are always subject to them. Being able to communicate depends upon this: knowing how to make communicative utterances, and knowing how to understand them depends upon knowing these constitutive rules.

The speaker's rule

The speaker's rule is that her utterance ought to be an attempt to get the hearer to recognise her as having the i/b/d/a that she actually does have. It is not simply that the speaker ought to be sincere, because sincerity implies a harmony between the i/d/b/a that the speaker has, and the one that *is represented in the utterance*, and yet this is not a requirement for successful communication: A sarcastic utterance, a joke, is not sincere in this sense, and yet the communication is successful if there is a harmony between what the speaker wants the hearer to recognise, and what the hearer does recognise from the utterance. A speaker can say 'Wow, this is a great party!' in a certain flat tone of voice in an attempt to get the hearer to recognise the fact that she thinks it is a terrible party, and this

does not flout the rule. [21] It is still conforming to the rule as long as the speaker wants the hearer to recognise her *actual* attitude.

Furthermore, non-conventional 'utterances' like gestures and noises might not easily fit into this picture of sincerity. A grunt, a shrug, raised eyebrows might not seem to be the kind of things that can be straightforwardly sincere or insincere.[22] However, for me, the constitutive rule only stipulates that the speaker wants the hearer to recognise her actual state of mind via recognition of her 'utterance', and thus raised eyebrows can obey or flout the rule. For example, if I raise my eyebrows in order to communicate to my friend that I am sceptical about an overheard claim, then I obey the rule if I am sceptical, and flout it if I am not.

So it is not a constitutive rule of successful communication that the speaker should be straightforwardly sincere in her utterance, but rather that there is a harmony between her actual i/d/b/a and the i/d/b/a that she intends the hearer to take her as having: she must be sincere *in her intentions*, rather than her utterance must be a sincere representation of her i/d/b/a. She must be trying to get the hearer to recognise her as having the i/d/b/a that she actually does have. Breaking this rule occurs when the speaker tries to get the hearer to think she has an i/d/b/a that she does not have. There has to be an *openness* to the speaker's communicative intentions, such that when the hearer understands her utterance, if he were able to directly 'see' her mental state, he would not be surprised.

A French teacher may use the phrase 'Quelle heure est-il?' in a French conversation class without wanting to know the time, and yet not break the rule. As long as the teacher is not intending for the student to think that she wants to know the time, there is no insincerity in intentions here. She is not trying to get the hearer to think that she has an i/d/b/a that she does not have.[23] What this highlights is that the conventional or non-

[21] Sarcasm, it is plausible to say, does break a rule of language. Perhaps one ought to say 'This is a great party' only if one thinks it is a great party. However, the rule it breaks will be a rule governing language use, and not a constitutive rule of communication. (Though the two will turn out to be connected, I'm sure, but there is no space to go into that here.)

[22] See Green (2007) for interesting discussions on how one can express one's internal states.

[23] One might ask, what exactly is the French teacher communicating to the student, if it is not 'What is the time?'? What the teacher is communicating is something like: her desire that the student answer the question in French to demonstrate his understanding and pronunciation - this does not entail a desire to know the time. The teacher and the student are both aware that they are in a class, the point of which is to teach the student how to speak French. This shared knowledge explains how it is that the teacher's utterance is understood by the student to be a request for an answer, but not a request for the time. It would not be a problem if the student told her the wrong time, but it would be a problem if he told her the time in English. There must be a

conventional utterance itself is not the vehicle for the constitutive rule for communication. Conventional language is usually the means by which communication occurs, and in order to function as said means, there will be rules that govern its use.[24] However, conventional language is not always necessary for communication to occur (though, of course, it often is, if what is to be communicated is more complicated than can be expressed with gestures and non-conventional noises), communication can successfully occur between two people who do not share a language.[25] The constitutive rules for communication governs *all* instances of communication, whether they are conventional or not. The rule for the speaker is that she only try to get the hearer to recognise her as having an i/d/b/a that she actually does have.

The necessity of this rule can be made apparent if we try to explain how communication might work if it was not the rule. If there was no rule such that the speaker ought to try and get the hearer to recognise her as having an i/d/b/a that she actually has, then there is no necessary link between her utterance and her mental state. If a speaker's utterances are not guided by this rule, then what is there to guide her communicative behaviour? If there were no such rule for the speaker, then there would be no way of distinguishing between her utterances that were attempts to get the hearer to recognise her actual mental state, and utterances that were in no way connected to her mental state. This is not to say that she would be sometimes lying, and sometimes not, because the distinction between lying and being 'sincere' depends upon the rule that utterances ought not to be attempts to get the hearer to think she has an i/b/d/a that she does not have. Without this rule, her utterances are equally valid, equally correct, equally 'sincere'.

---

common ground between speaker and hearer such that the speaker can reasonably expect the hearer to understand her utterance (her intentions) correctly. (Perhaps something like Schiffer's 'mutual knowledge' condition on communication [1972 p.30.])

[24] These rules may even be constitutive rules: knowing *how* to use language is synonymous with knowing the rules that govern it. (See Hanfling (1980) for discussion on whether language needs rules.) And yet these rules are rules for language use, not the rules for communication. I believe, though I will not argue, that the constitutive rules of communication are more fundamental that the rules of language use.

[25] It is a necessary condition on trying to communicate one's mental state, or on recognising the mental state of another, that there is a something like a shared language or common ground. (See n.23). The speaker, in attempting to get the hearer to think she has a certain mental state, must be reasonably sure that the hearer will be able to recognise her as having it and this means she must presume that her communicative act, her utterance, will be recognisable/understandable/meaningful to him. Her having this communicative intention entails that she makes this presumption about the hearer. And the hearer's ability to recognise the speaker's mental state upon hearing the utterance entails that he is able to understand what it means. However, language is just one of the means by which one can communicate and I hope that because language is not a necessary condition for communication, I can avoid getting into the (massive) debate on linguistic meaning; my focus is not on meaning, nor the origins of meaning, but on the possibility of communication. (Though it is foreseeable that my account could have ramifications in that area, it will not be helpful to go there now.)

The constitutive rule for the speaker explains the *necessary connection* between the speaker's mental state, and her utterance, without resulting in the unpalatable consequence that all utterances necessarily *are* attempts to communicate the speaker's actual mental state. Some are not, but they are normatively distinct from those that are. The normative distinction is essential if we are to understand the speaker's behaviour as communicative. If there is nothing to distinguish between utterances that are attempts to communicate a speaker's actual mental state, and utterances that are attempts to communicate something other than the speaker's actual mental state, then how can utterances be *attempts to communicate* a mental state? How can we make sense of an utterance being a means to getting a hearer to recognise the speaker's mental state? Without this rule, it becomes hard to understand a speaker's utterances as communicative behaviour at all.[26]

Furthermore, if the speaker's utterances are not guided by this rule, then it is hard to see how the hearer could have access to the speaker's mental state upon hearing an utterance. The hearer, upon hearing an utterance, would not have access to the speaker's mental state because there is no rule guiding her behaviour such that her utterances ought to be attempts to get him to recognise her actual i/b/d/a. Her utterances can be understood to be correct whether or not they are attempts to get him to recognise her actual mental state, and therefore he is in no position to recognise her actual mental state upon hearing her utterance. For example, a speaker could make the utterance 'I'm tired,' but if there is no rule such that she should only say 'I'm tired' (or mime that she is tired or whatever it is that she does in order to get the hearer to think she is tired) when she actually is tired, then it is unclear how her utterance could be a means to communicating to the hearer that she is tired. This rule needs to be in place in order to explain how it is that the hearer can come to recognise that she is tired upon understanding her utterance.[27]

[26] This fact, that it becomes hard to recognise her behaviour as communicative, highlights the fact that the rule is *constitutive* of communication, not just regulative. A non-constitutive rule is a rule that governs actions within an already existing kind of behaviour, e.g. good table manners. Table manners are rules that guide one's eating behaviour. But the behaviour of eating is something that exists with or without the rules. Communication as an institution, on the other hand, could not exist without the rules. The constitutive rules create the behaviour in question.

[27] This point is developed in Section Two, iii.

The hearer's rule

The hearer's rule, more than the speaker's rule, sets my account apart from most accounts of communication. The idea that *making* utterances is a rule governed activity (for example, Williamson's claim that making an assertion is governed by the 'knowledge rule'[28]) is not particularly controversial. The idea that the hearer's *understanding* of an utterance is also rule governed is more so. I, unlike most theorists, believe that the interpretation of utterances must be a rule governed activity. The hearer simply knowing the speaker's rule is not sufficient (on its own) to explain how he can come to recognise the speaker's actual mental state upon understanding her utterance. Merely knowing the speaker's rule would mean knowing what the speaker *ought* to be doing with her utterance, but does not tell him what the speaker is *actually* doing with her utterance. It should not be a requirement that the hearer, in order to properly understand an utterance, has to *infer* that the speaker is following her rule, he must be able to *recognise her as following it* upon understanding her utterance. This means there must be a *correct* way to understand her utterance, and this correct way involves recognising her utterance as an attempt to get him to recognise her actual mental state.

The hearer, upon hearing the utterance of the speaker, must be in a position to recognise, thereby, the speaker's i/d/b/a. The hearer must be in a position to understand the speaker *as agreeing with him* when she nods, or understand her *as wanting to know the time* when she asks 'What is the time?' or as wanting him to close the window when she asks 'Please close the window.' The hearer must be able to recognise the speaker's mental state simply by understanding her utterance. What is it that puts the hearer in a position to recognise the speaker's i/d/b/a upon understanding the utterance? He is in this position because he knows how to communicate. This means that he knows that he *must recognise the speaker as trying to get him to recognise her as having the i/b/d/a that she actually does have*. The speaker's nod or question or request gives the hearer access to the speaker's i/d/b/a because he can recognise her utterance as an attempt to get him to recognise her as having the i/b/d/a that she actually does have. Properly or correctly understanding the

[28] Williamson 2000, Chapter 11.

utterance *entails a recognition of the speaker's actual mental state*. Knowing how to communicate means knowing (though is not necessarily aware) that one's interpretation of communicative utterances is governed by this rule.

The essence of this picture is that when the hearer understands an utterance, he is, at the same time, recognising the speaker's actual mental state (when all goes well). His understanding the utterance does not leave it up for grabs whether or not the speaker has the i/b/d/a that she is trying to get him to recognise her as having. Another way of saying this is that understanding the utterance and recognising the speaker's actual mental state is one cognitive event: the understanding of the utterance *is* the recognising of the speaker's actual mental state. When these two come apart, something has gone wrong. When the hearer does not recognise the speaker's actual mental state (the one she is trying to get him to recognise her as having) he is not obeying his rule.

Again, this is not straightforwardly that the hearer ought to take the speaker as being sincere. Firstly, sincerity implies a necessary connection with the form (or *linguistic* content) of the utterance, and he ought not to take sarcasm or irony as sincere in this technical sense. Secondly, simply taking her as sincere does not necessarily mean that he would recognise her actual mental state: she could be misleading him; he could misunderstand her. In both these cases, the hearer will *think that* he is recognising the speaker's actual mental state, and yet not be. What he needs to recognise upon understanding the utterance is the speaker's actual mental state that she is trying to communicate to him. He needs to see her as intending to communicate to him the i/b/d/a that she *actually has*. Understanding properly requires that he would not be surprised if he were able to 'see' her mental state (i/d/b/a) directly. Understanding properly requires that he have access to her actual mental state. The constitutive rule for the hearer, the necessary rule, is that he take the speaker to be trying to communicate to him her actual mental state; when this rule is obeyed, the hearer has access to the speaker's actual mental state. It is this rule that allows for the possibility of communication, because it is what explains the speaker's communicative intentions, and explains how these intentions are fulfilled.

In order for the speaker to intend to communicate to the hearer a certain i/d/b/a through her utterance, she must think that the hearer is in a position to recognise her as having this i/d/b/a upon hearing and understanding her utterance. This means that when the hearer hears her utterance, he must be able to recognise her as having the i/b/d/a that

she is trying to get him to recognise her as having. It cannot be left open upon understanding the speaker's utterance whether or not the speaker has the i/d/b/a that she is trying to get him to think that she has. In order for the speaker to successfully communicate her i/d/b/a to the hearer, he must take her as having that i/d/b/a, as having the i/b/d/a that she actually does have. If her nod did not communicate to him that she agreed with him then she would not have succeeded in getting him to recognise her as agreeing with him, and yet her intention behind the nod was to do just that.

The necessity of this rule (that the hearer take the speaker to be trying to communicate an i/d/b/a that she *actually does have*) is apparent if we try to explain the speaker's communicative intentions in its absence. If the hearer is not bound to recognise the utterance as an attempt to get him to recognise the speaker as having the i/b/d/a that she actually does have, then he is in no position to recognise her as having that i/d/b/a upon hearing her utterance. He may recognise that she *wants him to think that* she has a certain i/d/b/a, but this is not enough to get him to realise that she actually does have it. Her nod does not allow him to see her as agreeing with him, her question 'What is the time?' does not allow him or put him in a position to recognise her as wanting to know the time. And if this is the case, how can it be the speaker's intention to get him to recognise her as having these i/d/b/a's upon hearing her utterance? The speaker's intentions are inexplicable if the hearer is not bound by the constitutive rule of communication to take her as having the i/d/b/a that she is trying to get him to recognise her as having.[29]

In order for the speaker to have the intention to communicate her i/d/b/a through an utterance, she would need to know or at least believe that the hearer is in a position to recognise her i/d/b/a via his *understanding* the utterance. She must think that he will be able to understand the language or gestures she employs and thereby come to recognise her as actually having the i/d/b/a that she is intending to communicate to him. She must think that her nod will get him to recognise her as agreeing with him, or her utterance 'What's the time?' will get him to recognise her as wanting to know the time, or that her utterance 'Please close the window' will get him to recognise her desire that he close the window. The form (or linguistic content) of the utterance itself is not central to her

[29] This idea is developed in more detail in Section Two, ii.

intentions, what is central is that the utterance (whatever it is) allows the hearer to see or recognise her as having the i/d/b/a that she is trying to get him to recognise her as having.

To be clear, the speaker's intentions cannot be explained by her knowledge that the hearer understands English, or is familiar with body language. The hearer's understanding the linguistic content of the utterance – knowing what it 'means' – though necessary, does not on its own, give him access to the speaker's mental state. What explains her intentions is that she knows that the hearer is able to *recognise her as having* the i/b/d/a that she is trying to get him to recognise her as having, upon understanding her utterance. This knowledge is available to her because she knows how to communicate, which means that she knows that her utterance is guided by the rule that she ought only to try to get the hearer to recognise her as having the i/b/d/a that she actually does have, and that the hearer, upon hearing and understanding an utterance is in a position to recognise her as having the i/b/d/a that she is trying to get him to recognise her as having (because this is how he ought to understand it). The hearer's ability to understand English is necessary perhaps, but it does not explain the speaker's communicative intention. His understanding of the literal meaning of the words she employs does not give him access to her mental state. What gives him access to her mental state is his ability to communicate, not just his ability to speak English.

The hearer's rule is that he recognise the utterance as an attempt to get him to recognise the speaker's actual mental state. Recognising the speaker's actual mental state is a necessary part of successfully understanding the utterance. But it is not sufficient. It is not only necessary that he recognise the speaker's actual mental state, but that he recognise that *this* is what she is attempting to communicate to him through her utterance. Sometimes a speaker may be trying to get the hearer to recognise her mental state as being something other than it actually is, and the yet the hearer may see through this and recognise her actual mental state. In this situation, the hearer is recognising the speaker's actual mental state, but he is not obeying his rule (and neither is the speaker obeying hers), because he is not seeing her as attempting (understanding her utterance as an attempt) to get him to recognise her actual mental state.

The rule for the speaker is that she ought only to intend to communicate her actual mental states, and the rule for the hearer is that he ought to recognise her as communicating her actual mental states. That these rules are constitutive of successful

communication means that they are in place or in force in *every instance* of communication. The rules are known by both the speaker and the hearer.[30] The speaker not only knows that her utterance ought to be an attempt to get the hearer to recognise her actual mental state, but she knows that the hearer ought to understand it as such. The hearer not only knows that he ought to understand a speaker's utterance as an attempt to get him to recognise her actual mental state, he knows that her utterance ought to be an attempt to do just that. The speaker, knowing the hearer's rule (and knowing that he knows the rule), knows that her utterance is a means to communicating to the hearer her actual i/b/d/a. The correct way to make an utterance is as an attempt to get the hearer to recognise one's actual i/b/d/a, and the correct way to understand an utterance is as an attempt to get the hearer to recognise the speaker's actual i/b/d/a. It is the existence of these rules that explains how it is that a speaker can communicate her i/b/d/a to a hearer with an utterance.

It is only when both the rules are conformed to that *successful* or *correct* communication takes place, but even when the rules are flouted and broken, communication can still occur. The possibility of communication depends upon the existence of the rules, and on the fact that participants are *necessarily governed by* or subject to these rules, but not on the participants necessarily *obeying* the rules. Communication still takes place when the speaker has hidden intentions, and even if the hearer does not recognise the speaker as intending to communicate her actual mental state, communication can occur, but this kind of communication is *necessarily* deviant or parasitic on successful or correct communication.

Breaking the speaker's rule

Imagine there is an elderly woman on a train. She feels a cold draught from an open window and, rubbing her shoulders, she says to a young man 'Brrrr, please would you close that window?' The man gets up and closes the window. This is a successful or correct instance of communication because the woman obeys her rule as the speaker, she tries to get the

[30] To re-emphasise: knowing the rules does not mean being *aware of* the rules. Knowing the rules is synonymous with *knowing how* to behave as a participant in a certain institution. I know *how* to understand my son's plea 'Bikkie, peese!' because I know that his utterance is an attempt to get me to recognise his desire for a biscuit. And he knows this too. This is what explains his utterance.

hearer to recognise her desire that he close the window, and her actual desire is that he close the window. The man obeys his rule, he understands her utterance to be an attempt to get him to recognise the old woman's actual desire that he close the window (that the window be closed). The hearer's rule is what explains the speaker's intention: the woman knows that her utterance will be a means to communicating her actual desire because she knows that the hearer is in a position to understand her utterance as an attempt to get him to recognise her actual desire (he knows the rules); and the speaker's rule is what explains the hearer's ability to understand her utterance as an attempt to get him to recognise her actual desire (that is, recognise her actual desire).

This instance of communication is successful because there is an openness, or a mutual awareness of intentions and desires between the speaker and the hearer. The woman rubbing her shoulders and saying 'Brrr, please could you close the window' gets the hearer to recognise that she wants him to close the window because she is cold, she doesn't like the draught. All this is clear to the man, and to the woman. There are no hidden intentions, and the communicative act gets the hearer to recognise the speaker's actual i/b/d/a. Her communicative utterance is an attempt to get the hearer to recognise her actual desire, and this is what he recognises upon correctly understanding her utterance. The actual form or content of the utterance, whether it be a question or assertions (such as 'I am cold', 'I would prefer the window to be closed', 'You are a tall young man, you could reach that window.') or a collection of gestures, (e.g., arm rubbing, exaggerated shivering, pointing at the window, pleading smile, etc.) is not an essential part of the story. The important part is that whatever the utterance is, it gets the hearer to recognise what the speaker intends him to recognise, and he recognises it because he recognises her as attempting to get him to recognise this with her utterance.

The speaker can break her rule, and yet still communicate. Perhaps the old woman is not cold, and she does not want the man to close the window, she wants him to stand up so that she can take his seat, or because she hopes that when he turns away she can steal his newspaper. In this instance she uses the utterance 'Brrr, please would you close the window?' to communicate a desire that he close the window (that the window be closed), but that is not her actual desire. If he were able to see her mental state he would be surprised: there is a lack of openness in her intentions. Her utterance is an attempt to get him to recognise that she is cold, and would prefer the window to be closed, and yet this is

not her actual internal mental state. She gets him to think that she wants the window to be closed, but she does not desire this. This is an instance of flawed communication.

In one important sense this instance of communication still succeeds: the woman succeeds in her aim of communicating to the man *that* she desires that the window be closed, and yet it is incorrect (*technically* unsuccessful) because she breaks the rules – she does not actually desire the window to be closed. We know that she is breaking the rules because it is the hearer's presumption that she is *obeying* the rule that allows her to succeed in her hidden intentions. When the man hears the woman's utterance 'Brrr, please would you close the window?', he is in a position to 'recognise' her as being cold and wanting him to close the window (wanting the window to be closed), and he is in this position because he knows the rules of communication. If he did not know this, if he did not know that the speaker's utterance is meant to be understood as an attempt to get him to recognise her actual desire, then he would not be in a position to understand her utterance "Brrrr, please would you close the window?' as an attempt to get him to recognise her actual desire. It is only his presumption that her intentions are open to him that allows her intentions that are *not* open to him to be fulfilled.

At this point one might claim that there is no hidden intentions here, the woman wants the man to close the window in the sincere instance because she is cold, and she wants him to close the window in the second instance for some other reason: but in both instances, one could argue, she *does want him* to close the window. However, it is not straight forwardly the case that she does want him to close the window in the second, deceptive, case – she does not want the window to be closed. The distinction is clear if we see that if the woman made her *actual* mental state explicit in both cases, the hearer would recognise different things. In the instance where she is being sincere, he recognises that she wants the window closed, *and* that she wants him to close it. In the insincere case, he would not recognise that she wants the window closed, only that she wants him to close it. This is not the same thing. Her actually wanting the window to be closed is an important part of what she wants him to recognise when she asks him to close the window. It is this that is meant to motivate him to close it: her wanting it closed would be his reason for closing the window. It is only because he takes her utterance to be an attempt to get him to recognise her *actual* mental state of wanting the window to be closed, that she can possibly succeed in her devious designs. If he were to recognise her actual mental state, her actual

desire, (that she wants him to get up and leave his newspaper unattended) she could not expect her utterance to have the (extra-communicative) effect she intends.

The speaker breaking her rule is an instance of incorrect communication, technically unsuccessful. It might seem intuitively strange to claim this, given that the speaker's communicative intentions are fulfilled: she intends to get the hearer to think that she has a certain i/b/d/a, and succeeds in doing this (as long as the hearer obeys (or thinks he is obeying) his rule.) And yet, it *is* incorrect, and in the sense that successful communication entails the communication of the speaker's actual mental state, it is unsuccessful. The fact that the hearer also necessarily breaks his rule when this happens, is testament to this fact. The hearer, upon hearing an utterance ought to understand it as an attempt to get him to recognise the speaker's actual mental state, and if she is deceiving him then he does not do this. Even if the speaker's intentions are fulfilled, their being hidden and not open, the fact that the hearer would be surprised if he could 'see' her mental state directly, means that there is something wrong with this instance of communication.

The flawed kind of communication that occurs when the speaker breaks her rule is *parasitic* on successful communication. This is because it is only possible for the speaker to get the hearer to think she has a certain mental state because she knows that the hearer ought to think she actually has the mental state that she is trying to get him to think that she has. The existence of the rules is what allows for the possibility of using utterances incorrectly to communicate with a hearer. If it were not a rule that the hearer take an utterance as an attempt to get him to recognise the speaker's actual mental state, then when the speaker uses an utterance incorrectly, the hearer would not understand her utterance as an attempt to get him to recognise her actual mental state. The woman who uses her utterance to try to get the hearer to *think that* she wants him to close the window is able to do so only because the hearer ought to understand her utterance as an attempt to get him to recognise her as *actually wanting* him to close the window. The possibility of unsuccessful communication occurring depends upon the existence of the rules as much as the possibility of successful communication.

Breaking the hearer's rule

The hearer's rule is that he recognise the speaker as trying to get him to recognise her as having the intention/desire/belief/attitude that she actually does have, and this rule can be broken in two ways. It is possible for the hearer to hear an utterance and reject it as the speaker trying to get him to recognise her actual mental state, and it is possible for the hearer to *think that* he is recognising the speaker as trying to get him to recognise her as having the intention/desire/belief/attitude that she actually does have and yet not be. To break the rule in the first way is to reject the idea that the speaker is being sincere in her intentions. Let's say the young man, for example, has seen this particular elderly woman around before, and is familiar with her trick of getting people to move so that she can take their newspaper. In this instance, when she asks him to close the window, he does not understand her utterance to be an attempt to get him to recognise her actual desire that he close the window. In this instance, the hearer breaks his rule, and it is not clear whether or not communication has occurred. Unlike when the speaker breaks her rule, when the hearer breaks his in this way, he *blocks* the communication, he does not accept the utterance in the way that the speaker intends him to accept it.

I am not sure whether or not to say that communication (faulty or parasitic) occurs here. Depending on one's intuitions, one could say that the woman has communicated *that she wants him to think* that she wants him to close the window. It is clearly not her intention to communicate this to him, but this may not be fatal to the idea of communication occurring in a broad sense. However, whether one wants to claim that communication occurs in this instance or not, it is clear that *if* it does, it is not *correct* communication. Whatever happens, it is not what the speaker intended to happen when she made her utterance, and it is difficult to apply the word success to such an occurrence. The hearer breaks his rule by rejecting the speaker's utterance as an attempt to get him to recognise her actual mental state: communication technically fails in this instance.

The second way for the hearer to break his rule is to fail to recognise the speaker's actual i/b/d/a. This can happen either because the speaker is breaking her rule (she is not trying to get him to recognise her actual i/b/d/a) or because he *misunderstands* her utterance. If the woman from the above example succeeds in getting the man to think that she wants him to

close the window when this is not, in fact, her actual desire, then he thinks that he has successfully recognised her i/b/d/a, but he has not. What this means is that if the speaker breaks her rule, the hearer will necessarily break his, either by rejecting her utterance, or by taking it to be sincere when it is not.[31] Correct understanding entails recognising the speaker as trying to get the hearer to recognise her actual mental state, and this cannot happen if the speaker is not trying to get the hearer to recognise her actual mental state.

But even if the speaker obeys her rule, it is possible for the hearer to think he is recognising her actual mental state, and yet not be. There can be a *mis*understanding. Say, for example, the woman asks the man 'Please would you close the window?' and yet the noise of the wind and the train drowns out her utterance, or the young man is not listening properly. He sees that she is requesting something, but believes it to be something other than that he close the window. Perhaps he thinks she says 'Please may I borrow your newspaper.' In this instance, the speaker is not breaking her rule, and the hearer is attempting to follow his rule, and yet he fails to do so. The hearer's rule is that he ought to recognise the speaker as trying to get him to recognise her as having the intention/desire/belief/attitude that she actually does have, and obeying this rule entails recognising her actual mental state. If he does not recognise her actual mental state, then he fails to conform to the rule.

Perhaps a more common example of misunderstanding would be a case where the speaker uses a non-conventional gesture to communicate a certain i/b/d/a. Perhaps a woman raises her eyebrows to communicate to her friend that she thinks it is time to go. The friend, seeing the raised eyebrows, believes that she is trying to communicate to him that she is sceptical about a comment overheard. In this case the speaker obeys her rule; the hearer, thinking he is obeying his rule, takes her to have a certain i/b/d/a that she in fact does *not* have. Communication does occur in this instance, and yet it is a failed or deviant kind of communication: something has gone wrong, there is a *mis*understanding. In this instance, what has gone wrong is that the hearer has not conformed to the rule such that he

[31] This, at first sight, looks like a straight forward case of an ought where there is no 'can'. A hearer ought to understand an utterance as an attempt to get him to recognise the speaker's actual mental state, but if the speaker breaks her rule, he cannot do this. However, this would imply a very strict understanding of Kant's principle such that it must *always* be possible to do what one ought. But it must be allowable that sometimes circumstances prevent one from doing what one ought to do.

ought to recognise the speaker as trying to get him to recognise her actual mental state. He can fail to obey the rule even if he is attempting to obey it.[32]

When the speaker breaks her rule, she usually does it intentionally, with a certain goal in mind.[33] The act of trying to get a hearer to think one has an i/b/d/a that one does not have is an intentional action. When the hearer breaks his rule, on the other hand, there is no intention to do so: there is no goal in mind. If he rejects the utterance, it will usually (though not always) be because he has some evidence to think that the speaker is breaking her rule: there will be a 'defeater', something that defeats the hearer's rule that he take her utterance as an attempt to get him to recognise her actual mental state. It would not seem accurate to describe his rejection as an intentional breaking of the rules. And if he accepts her utterance, and yet does not recognise her actual mental state (because she is being insincere), he does not even realise that he is breaking the rule, never mind has an intention to do so. Because of this, it might sound strange to say that the hearer is *breaking a rule*, when his action is not an intentional action in the way that the speaker's breaking the rule is.

But this is not problematic if one considers the nature of constitutive rules, and the normative force that they exercise. Knowing the rules is synonymous with knowing how to communicate. The rule that a hearer ought to understand a speaker as trying to get him to recognise her as having an i/b/d/a that she actually does have is not a rule that one has a moral or epistemic obligation to obey. It is simply a rule that necessarily governs the hearer's understanding of utterances: it is what puts him in a position to understand a communicative utterance. The rule guides his behaviour in that it tells him *how to* interpret communicative utterances. (In the same way that knowing *how to* play a game means that one's actions, whilst playing, are governed by the rules of the game.) Breaking the rule is not

[32] This issue of misunderstanding brings another possibility to the surface. Consider a situation where the woman wants the man to close the window, asks him to, but the wind is so loud he can't hear her. However, he still comes to believe that she wants him to close the window based on other clues, not her utterance. Would this count as successful communication? It fulfils some of the important criteria for successful communication, but, under my account, this is not technically successful. This is because the hearer's recognition of the speaker's mental state is not his *understanding of the utterance*. Successful communication on my account occurs when the hearer recognises the speaker's actual mental state upon his understanding of her utterance. However, possibilities such as these warrant better discussion than I can afford here.

[33] It should be noted that, when the utterance is an instance of *testimony*, the speaker can break her rule unintentionally, but that is the subject of the next chapter.

necessarily an intentional action, and can happen even when the hearer believes he is following it. For example, it is a rule of tennis that the ball must land inside certain lines, but the ball can land outside the lines without it being the intention of the players that it does, and it may even land outside the lines without the players being aware that it has.

The possibility of communication depends upon the existence of these rules. The speaker and the hearer must know how to communicate, and in knowing this they know that speakers ought to try to communicate their actual i/b/d/a, and that hearer's ought to understand them as so doing. This 'know-how' allows for communication to occur, and when the rules are obeyed it allows for communication to properly succeed. When both of the rules are obeyed, the speaker manages to communicate her actual i/b/d/a to the hearer. In making her utterance she makes her i/b/d/a available to the hearer, and in properly understanding that utterance, the hearer recognises her as having the i/b/d/a that she actually does have. The speaker's actual mental state is laid bare or revealed to the hearer upon a successful instance of communication: if it is not (if she is trying to get him to see her as having an i/b/d/a that she does not have, or if he does not recognise her as having the i/b/d/a that she is intending for him to recognise her as having, or if he recognises her as having an i/b/d/a that she does not have) something has gone wrong. If communication does occur, it is necessarily a deviant, flawed or parasitic kind of communication.

Section Two: The unsatisfactory nature of thin accounts of communication

In this section I compare the explanatory ability of my account of communication with that of thinner, more popular accounts of communication defended by speech act theorists and communication theorists.[34] The fundamental difference between my account and thinner accounts is that my account proposes that successful communication entails the hearer coming to recognise the speaker's *actual* mental state: successful communication is *factive.* Other accounts of communication deny that there is a necessary link between the speaker's actual mental state and successful communication. Thinner accounts allow (in different

[34] These include, amongst others, Searle's speech act theory, Bach and Harnish's account of linguistic communication (1979), Strawson (1964), Brandom (1994), J.L. Austin (1962).

ways) for communication to count as a success even when the hearer does not come to recognise the speaker's actual mental state. I believe that this is a failing, and that it is a failing the consequences of which have not been fully appreciated. The aim of this section is to defend this claim, by explicating just how the lack of factiveness in thinner accounts undermines their explanatory powers.

In my view, a good account of communication should be able to explain:

i. How a speaker can communicate her internal mental states to a hearer,
ii. the existence of the speaker's communicative intention when making an utterance,
iii. and how an utterance comes to *mean* something to the hearer.

I will argue that the thinner accounts of communication fail to explain any of the above satisfactorily.

The picture that I put forward in the previous section is an attempt to get to the bare bones of communication, and what I hope I have revealed is that successful communication entails that what is communicated is *factive*. What I mean by this is that the speaker's *actual* i/b/d/a has a role to play in explaining how communication can occur. This is in direct contrast to the following encyclopaedia entry on speech acts by Kent Bach:

> an utterance can succeed as an act of communication even if the speaker does not possess the attitude he is expressing: communication is one thing, sincerity another. Communicating is as it were just putting an attitude on the table; sincerity is actually possessing the attitude one is expressing. Correlatively, the hearer can understand the utterance without regarding it as sincere, e.g., take it as an apology, as expressing regret for something, without believing that the speaker regrets having done the deed in question.[35]

Bach's thin conception of communication detaches the speaker's actual mental state from the process. It is not necessary that the speaker have the i/b/d/a that she is trying to communicate to the hearer that she has, and it is not necessary that the hearer take the speaker to have the i/b/d/a that she actually does have. Under this conception communication succeeds when the speaker manages to get the hearer to recognise her mental state, but equally succeeds when the speaker fails to get the hearer to recognise her as having the mental state that she wants him to think that she has. And it still counts as a

[35] Routledge Encyclopaedia.

successful instance of communication when the hearer thinks that the speaker has an i/b/d/a that she does not have. Under this conception, there is no normative distinction to be made between these instances of communication.[36] They are all equally communicatively successful.

For speech act theorists like Bach, the utterance puts the hearer in a position to recognise that the speaker has 'expressed' a certain attitude. By expressed, he means: 'For S to *express* an attitude is for S to R-intend [reflexively intend] the hearer to take S's utterance as a reason to think S has that attitude.'[37] Successful communication occurs when the hearer recognises that the speaker has expressed an attitude in this (non-factive) sense. For Searle 'In the case of illocutionary [speech] acts we succeed in doing what we are trying to do by getting our audience to recognize what we are trying to do. But the 'effect' on the hearer is not a belief or response, it consists simply in the hearer understanding the utterance of the speaker.'[38] Understanding the utterance consists in recognising the speech act that is being performed, or recognising the attitude that is being expressed, and does not entail recognising the speaker's actual attitude or mental state. So, what follows is a defence of the claim that such a conception of communication fails.

i. Under thin accounts of communication a hearer does not have to recognise the speaker's actual mental state (i/b/d/a) in order for communication to succeed. So, by definition, communication does not explain how a hearer comes to recognise the speaker's actual mental state. This is considered acceptable, as a hearer actually coming to think that the speaker has the mental state she is expressing is considered to be 'more than just communication.'[39] Communication is not the communication of mental states, it merely entails the speaker performing a certain speech act, and the hearer recognising that act for what it is. The speaker's actual mental state is not a part of the story.

Under my account, in contrast, successful communication entails the communication of the speaker's actual mental state. So, for example, if I want some ice-cream, attempt to get you to recognise my desire for ice-cream (with an utterance such as 'Please can I have

[36] There is often a sincerity rule posited, but this rule is not seen to be constitutive of communicative success, so it does not play a role in the explanation of how communication works. Utterances being sincere, and being taken as sincere, is something that happens *over and above* the instance of communication itself.
[37] 1979 p.15.
[38] 1970 p.47.
[39] Bach and Harnish 1980 p.16.

some ice-cream' or a nod in the direction of the ice-cream stand), and you successfully recognise my desire for ice-cream, then communication has succeeded. Both I (the speaker) and you (the hearer) have successfully conformed to the constitutive rules of communication: I have succeeded in communicating to you that I desire ice-cream. If I break my rule, such that I do not actually want ice-cream, even though I attempt to get you to think that I do, or if you break your rule, either by not recognising that I want ice-cream, or thinking that I do when I actually do not, then something has necessarily gone wrong with that instance of communication. The rules explain how the speaker communicates her mental state to the hearer because the hearer knowing the rules (knowing how to communicate) puts him in a position to understand utterances correctly.

The thin accounts can simply shrug at this point, for they have already conceded that communication does not require that the hearer recognise the speaker's mental state, and thus, my criticism has no sting. But, I contend, it should. The thin view of communication leaves what we would intuitively call 'successful communication' unexplained. In my view, it is counterintuitive to suppose that I have successfully communicated with you when I make the utterance 'Broccoli is disgusting' and you do not come to believe anything about my opinion on broccoli. It is counterintuitive to suppose that I have successfully communicated with you when your understanding of my suggestion 'Let's get some ice-cream' leaves it open whether or not I would like to get some ice-cream. It is counterintuitive to suppose that when and *if* you ever come to think that I would like some ice-cream, I have not *technically* communicated such to you. All I have communicated to you (under thin accounts) is that I have expressed an attitude, your taking me to have that attitude is something other than communication.

The thin accounts still want to maintain that there is *something* about the communicative act that explains how a hearer comes to recognise the speaker's mental state: the communicative act, according to Bach, provides the hearer with a *reason* to think that the speaker has the attitude expressed. But, on what grounds is the communicative act a reason to take the speaker as having the attitude expressed? As has already been established, the utterance is as communicatively correct when it is insincere as when it is sincere, so, in actual fact, the utterance says nothing about the speaker's actual mental state. The question is, how can the communicative act (thinly conceived) count as a reason for the hearer to take the speaker as having a certain attitude or mental state? There is

nothing binding the speaker to be sincere, and nothing binding the hearer to take an utterance as sincere. It is left wide open, under thin accounts, how it is that an utterance can get the hearer to recognise the speaker's mental state.

If we are generous, and allow that the utterance does (somehow) provide the hearer with a reason for taking the speaker to have the attitude she is expressing, this is still not a satisfactory explanation of how the hearer comes to have access to the speaker's mental state. This is because, the hearer has this reason to take the speaker as having the i/b/d/a in question whether or not the speaker actually has the attitude. So, if all goes well, the speaker makes her utterance (expresses an attitude), the hearer understands it, takes it as a reason to think the speaker has the attitude in question, and accepts that the speaker has the attitude, this does not explain how the hearer *has access to* the speaker's attitude. He has access to her utterance, which does not give him access to her mental state, and when he presumes that she has a certain mental state, it is still *a presumption*. He does not *know* what her mental state is. It is still an *inference* from hearing an utterance that the speaker has the mental state expressed, and it is an inference based on hearing an utterance that says nothing about the speaker's actual mental state.

Searle, and others[40], posit a sincerity rule, such that utterances ought to be sincere reflections of the speaker's mental state. If this is a rule, perhaps the hearer is thereby put in a position to recognise a speaker' utterance as sincere. But the hearer recognising an utterance for what it is, understanding it properly, does not put him in a position to recognise the speaker's actual mental state. Making a (non-constitutive) sincerity rule comes too late. What puts the hearer in a position to know that the sincerity rule has been conformed to? There is nothing about the utterance itself that reveals its sincerity to the hearer, so he would have to make some kind of *leap of faith* that the speaker is being sincere in order to think that the speaker has the mental state she is trying to get him to think she has. He has to presume that she is conforming to the rule. And yet even if he does presume her to be being sincere, and she does happen to be being sincere in this instance, it seems to me that he will *never know* what her mental state is, it will always be a matter of presumption.

[40] e.g. Austin proposes that sincerity is a 'felicity' condition on speech acts (1962).

If, as on my account, the 'sincerity rule' is constitutive of successful communication (utterances ought to be attempts to get the hearer to recognise the speaker's *actual* mental state), and the hearer's rule is such that he ought to recognise it as such, the problem does not arise. This is because the hearer, upon properly understanding an utterance, is not required to presume that the speaker is being sincere, he does not need to make a leap of faith, or *infer* that the speaker has the mental state she is expressing. Upon properly understanding the utterance, the hearer recognises the speaker's actual mental state. Successful communication entails recognition of the speaker's actual mental state. Understanding the utterance leaves no gap between recognising the mental state being expressed and coming to see that the speaker has that mental state. Understanding the utterance *is* coming to see the speaker as having that mental state. If he does not, then something, necessarily, has gone wrong. For speech act theory, how the hearer comes to recognise the speaker's actual mental state is (I contend) an epistemic mystery.

In my view this is a failing that the communication theorists should take seriously. However, as it is commonly accepted that communication by itself is not the means by which a speaker communicates her mental state to the hearer, it is technically permissible for communication theorists and speech act theorists to ignore this objection. They could be justified in claiming that the problem of explaining how a speaker gets a hearer to recognise her actual mental state is not a problem for communication theorists. The problem is, however, that the thinness of their account is not only a problem of intuitive satisfactoriness of their account, it is a problem when it comes to making sense of their own presumptions.

ii. Under my account, and under thin accounts, the speaker intends for her utterance to be a means to 'communicating' her mental state to the hearer. When I say 'Let's get some ice-cream' I intend for you to recognise my desire for getting ice-cream. When I say 'Broccoli is disgusting' I intend for you to recognise my opinion that broccoli is disgusting. This is my (what I will call) *communicative* intention when I make these utterances. That there is such an intention is accepted under the thin account of communication too: for Bach, the speaker intends for the hearer to take the utterance *as a reason* to think that she (the speaker) has the attitude expressed. That the hearer comes to think the speaker has the attitude expressed is part of the intention behind making an utterance. For Searle, speech

acts are identified by their essential rule, which, for questions is that they 'count as an attempt to elicit information from H,'[41] or for requests, that they 'count as an attempt to get H to do A.'[42] In order for the utterances to 'count as' the speech acts that they do, they would have to be a means to communicating to the hearer that the speaker wants the information, or wants H to do A.

The speaker's intention to get the hearer to recognise the speaker as having a certain mental state is recognised by thin accounts of communication as part of the speaker's intentions when performing a speech act, but they would not call it *communicative* intention because communication, for them, does *not* entail the hearer coming to recognise the speaker as having the mental state she is expressing. For speech act theory, the intention that does the communicative work is the 'illocutionary' intention[43], and this intention is the same whether or not the speaker has the attitude she is expressing. The illocutionary intention is that the hearer recognise what the speaker is doing with her speech act. The illocutionary intention is the intention that is fulfilled once the hearer has recognised *what attitude is being expressed*, and not when he recognises the speaker *as having* a certain attitude. The intention that the hearer actually *take the speaker as having* the attitude she is expressing, on the other hand, is the intention that is *not* fulfilled, on thin accounts, when communication is successful. This communicative intention, however, is characteristic of all instances of communication; the point of performing a speech act, or of making a communicative utterance to a hearer, is to get the hearer to recognise the speaker as having a certain mental state.[44] If the speaker asks a question, she wants the hearer to think that she wants some information, if she expresses an opinion, she wants the hearer to think that this is her opinion.

---

[41] 1979 p.66.
[42] *Ibid.*
[43] First identified by Austin (1962).
[44] Speech acts lacking this intention could not accurately be described as communicative acts, they would be acts that have no communicative purpose. The communicative intention, as I have defined it, is not the perlocutionary intention. There seems to be some ambiguity within speech act theory: on the one hand, the illocutionary intention is fulfilled even if the hearer does not take the speaker to be being sincere, and yet the perlocutionary intention of the utterance is that the hearer *act on* or *respond to* the utterance. The level at which the hearer *accepts* the utterance as sincere (takes the speaker to have the attitude she is expressing) is not clearly defined – this is probably because very little attention is paid to the role of hearer understanding. (Searle, at first, makes the claim that the hearer accepts an utterance (a promise) as sincere upon hearing it, but later amends this analysis to allow for insincere promises by claiming that the hearer doesn't need to recognise the speaker as being sincere, and thus, in my book, undermines his entire theory. 1990 pp.13-14.)

Under my account this intention is explicable, because a successful instance of communication entails that the hearer recognises, upon understanding her utterance, her mental state. This means that the speaker can intend for her utterance to be a means to communicating her mental state to the hearer, because, all going well, this is what *will* happen when she makes her utterance. Her utterance counts as an attempt to get the hearer to recognise her actual mental state because her action is governed by the constitutive rule such that her utterance ought to be an attempt to get the hearer to recognise her actual mental state. Furthermore, the hearer's understanding of the utterance is governed by the constitutive rule such that he ought to recognise her utterance as an attempt to get him to recognise her actual mental state. Both the speaker and the hearer know the rules of communication (they know how to communicate) and this explains the speaker's communicative intention when she makes an utterance.

The thin accounts of communication do not have the resources to explain the communicative intention that the hearer take the speaker to have the attitude she is expressing. For them, an instance of communication is successful if the hearer recognises what attitude the speaker is expressing. He does not need to take the speaker as having that attitude. And if this is the case, what is it that motivates or explains the speaker's intention to get him to think she has a certain attitude, if it is not a part of communicative success that he come to think she has that attitude? The utterance itself gives the hearer access to the fact that the speaker has performed a certain speech act, or that she has 'put a certain attitude on the table,'[45] but it does not tell him what her attitude is. If this is the case, why does the speaker make the utterance? How is it that she sees her utterance as a *means* to getting him to think she has a certain attitude if the utterance does not give the hearer access to her attitude?

The speaker's intention behind the utterance is, it seems to me, inexplicable from within the thin account of communication. However, one could argue that the utterance does give the hearer access to the speaker's mental state, because he has been presented with *a reason* to think she has that mental state. Perhaps, the fact that most communicative utterances are made sincerely means that the hearer can have access to the speaker's mental state, because, more often than not, when the speaker expresses an

[45] Bach, Routledge Encyclopedia entry.

attitude, she does have that attitude. The thin version of communication on its own does not offer an explanation for the speaker's 'communicative' intention (that the hearer take the speaker to have the attitude expressed) but this is not to say that other contextual or pragmatic considerations cannot give rise to the intention behind an utterance. The speaker may know that the utterance itself does not give the hearer access to her mental state (will not get him to take her as having the mental state expressed) but she may think that making the utterance will, at least, give him a reason to think she has the attitude nevertheless.

This may be so, and I do not have the space here to take on this claim. However, there are a few reasons for thinking that such reasoning will not suffice to explain the speaker's intention. Firstly, the idea that the speaker hopes the hearer will infer that she has the attitude expressed is not reflective of our experience of making utterances. When I say 'Let's have some ice-cream,' it seems strange to imagine that my intention is that you *work it out* from this that I want ice-cream. My saying it seems like a straightforward means to getting you to recognise my thought that we should get some ice-cream, and not, as the thin accounts contend, that I am offering you some kind of evidence, and you are using this evidence to conclude that I think we should get some ice-cream.[46]

Grice's account of non-natural meaning is summarised thus: '"A meant$_{NN}$ something by x" is (roughly) equivalent to "A intended the utterance of x to produce some effect in an audience by means of the recognition of this intention".[47] His account maintains the intuition that when a hearer understands an utterance (grasps the meaning of the utterance) the speaker's intended effect will be achieved. If the speaker's intention is that the hearer come to believe that p, then this is what occurs when the hearer recognises this intention. Speech act theorists since Grice have rejected this picture, claiming that the 'effect' produced is perlocutionary (external to the communicative act)[48] and that what the speaker must recognise is only what the speaker *means* by her utterance. However, what Grice's account captures that subsequent accounts of communication do not, is that there is an important sense in which the mere recognition of the speaker's intentions gives the

[46] There is a parallel here between this claim and the trust/assurance approach to testimony, where it is argued, correctly I believe, that asserting that p is plainly *not* a case of offering evidence for the fact that p. (See Ross 1986, Moran 2005 and Faulkner 2007.)
[47] 1957 p.385.
[48] Searle 1970 pp.45-50, Davis (1980).

hearer access to the speaker's attitude. It seems more straightforwardly the case that when I make an utterance I am *thereby* making my mental state available to the hearer, and not that I am hoping that he will use my utterance as a premise in an inference, the conclusion of which is that I have the attitude I am expressing.

A second reason to doubt that 'extra-communicative' considerations can explain the speaker's intention, is that, if the utterance itself does not give the hearer access to the speaker's mental state, other factors that could do the job are hard to come by. To see why, consider the seemingly helpful claim that 'most utterances reflect the speaker's actual mental state.' This claim would provide grounds for the hearer to take the speaker at her word (and thus explain the speaker's communicative intention). And yet, on what grounds can a hearer presume that most utterances are sincere, given that no utterance has given him access to the speaker's mental state? The hearer must always presume or infer the speaker's mental state from understanding the utterance. This means that in order to establish that most are sincere, he must presume that most of his presumptions were correct, and yet it is exactly his grounds for taking them to be sincere in the first place that are in question.

One could claim that if too many utterances are insincere, then there would be no such thing as communication: the very practice depends upon mostly sincere utterances.[49] However, although this might be true, it would not, on its own, explain the speaker's intention. The speaker would have to presume that the hearer knew this fact, and that from knowing it would infer that her particular utterance was sincere. Both presumptions are a little farfetched. The question here is not the one that epistemologists of testimony are at pains to respond to, namely *what justifies the hearer's testimony-based belief?* The question is why does the speaker make the communicative utterance with the intention of getting the hearer to think she has a certain mental state? This cannot be answered by explaining how the hearer is justified in taking her at her word. A satisfactory response must show why or how it is that the speaker sees her utterance *as a means* to getting the hearer to recognise her mental state.

[49] This is essentially C.A.J. Coady's argument against the Humean approach to testimony in his book *The Epistemology of Testimony*. The parallels here between the problems of communication and the problem of testimony are not coincidental. The epistemic gap to be crossed by the hearer in order to recognise the speaker's actual mental state is the gap that I deny exists in order to explain how testimony is a source of knowledge. See next chapter.

A third reason to doubt that extra-communicative considerations explain the speaker's intention, is that it is just not the case that the speaker considers such things when making an utterance. In order for it to explain (or rationalise) her behaviour, these reasons, whatever they may be, must play a role in her having the communicative intention that she does. In order to explain the speaker's intention satisfactorily, the thin conception of communication must show how the speaker comes to the conclusion (or knows) that the hearer will be in a position to work out what her attitude is, based on his understanding her as expressing a certain attitude. But making utterances does not seem to require the speaker having these reasons. Children communicate, and intend to communicate, with their utterances, without having any beliefs about how the hearer comes to recognise their attitudes. A much more satisfactory response to the problem is that the speaker's intention is fully explained by the possibility of communication itself – knowing how to communicate means knowing how one's utterances will be understood by the hearer (all going well).

iii. The third aspect of communication that thin accounts of communication fail to explain is how it is that an utterance can have 'illocutionary force', or come to have the 'meaning' that they do to a hearer. One of the central themes in speech act theory is that speech acts come in different types (e.g. requests, commands, assertions and questions), and that the type of action that an utterance is depends upon the *illocutionary* intention of the speaker – the utterance comes to have an *illocutionary force*. The hearer, upon recognising this illocutionary intention, recognises the speech act for the kind of act that it is, recognises it as a request, command, or assertion, etc. This is what constitutes (for speech act theorists) successful communication. However, although this is central to the theory, I contend that without making a *necessary connection* between the attitude expressed and the speaker's actual attitude, this idea (of illocutionary force) cannot get off the ground. In other words, it is my belief that the severing of the connection between utterance meaning and the speaker's actual mental state is fatal to speech act theory. My challenge is, if there is nothing essential that distinguishes a sincere utterance from an insincere one, if an utterance is just as communicatively correct whether or not it expresses the speaker's actual mental state, how does the utterance *count as* an expression of a certain mental state? If a speaker, in a particular instance, can say 'Let's get some ice-cream,' just as correctly as 'Let's get some hotdogs,' or even 'Opera is lovely', if any utterance is as correct

or communicatively successful *whatever the speaker's actual mental state*, then what is it about the first utterance that makes it a suggestion to get ice-cream? Why would the second and third utterances *not* be suggestions to get ice-cream?

There must be something that makes it the case that an utterance is an expression of a particular attitude, as opposed to another attitude. Under thin accounts of communication, it seems that there is nothing that can fulfil this role. If it is allowed that an utterance can be communicatively correct whether or not the speaker has the attitude in question, then what makes it an expression of *that* attitude as opposed to any other attitude? There may be an initial inclination to claim that the *meaning of the words* 'Let's get some ice-cream' has an essential connection to the attitude of wanting to get some ice-cream. But, and this is something speech act theorists propose, the illocutionary act, the act being performed, cannot be understood from merely understanding the 'locutionary act'. The locutionary act is the act of saying something, using words with a certain meaning. Speech act theory, and theory about communication, is interested in explaining how the speaker can use locutionary acts to perform illocutionary (communicative) acts. The illocutionary force of the utterance (what the hearer must recognise in order for communication to succeed) is supposed to be created *solely* by the speaker's illocutionary intention and *not* the meaning of the words used.

The words 'Let's get some ice-cream' do not count as a suggestion that one get ice-cream just on account of the meaning of the words. (They could be the lyrics to a song, part of a joke, etc.) The speaker has to be *using* them with the intention of making the suggestion. It is the speaker's illocutionary intention that *defines* the act for what it is. And this is the point at which I am making my complaint: how does the use of the words 'Let's get an ice-cream' come to count as, come to *be*, a suggestion that one get ice-cream if there is no necessary link between the utterance and the speaker's actual mental state. If the utterance is as communicatively successful whether or not the speaker wants to get ice-cream upon making the suggestion, or whether or not the hearer takes her as wanting to get ice-cream, then what is it about the utterance that makes it a suggestion to get ice-cream? Where does the illocutionary force come from if the speaker's actual mental state at the time of making the utterance is neither here nor there?

For Searle, there are 'counts as' rules (essential rules) that are in play for each kind of speech act.[50] So, for example, the constitutive (or essential) rule for a request is that it 'counts as an attempt to get H to do A.' and the essential rule for giving advice is that the utterance 'counts as an undertaking to the effect that A is in H's best interest.'[51] These essential rules are known by both speaker and hearer, and are what explain how the utterance can have the illocutionary force that it does. The speaker knows her request counts as a an attempt to get H to do A, and the hearer, upon understanding the utterance, understands that the speaker's utterance is a request, and thus counts as an attempt to get him to do A. However, my claim is that this essential rule as Searle states it, cannot function in the way that he intends unless he concedes that there is something about the utterance that links it to the speaker's actual mental state. For one thing, if there is nothing about the utterance itself that connects it to the speaker's actual mental state, then knowing the counts as rule does not give the hearer access to what the speaker is doing with the utterance. Knowing that requests 'count as' attempts to get H to do A does not put the hearer in a position to recognise an utterance as a request.

To put this another way, it is my belief that the 'counts as' rule is essentially circular. For Searle, it is on the basis of the constitutive rule for the different speech acts that the hearer can recognise what act is being performed. But on what grounds does he recognise that the utterance he has just heard is of the type that counts as certain speech act? In order to understand an utterance as a request, he must recognise that the utterance 'counts as an attempt to get H to do A,' and yet in order to recognise that the utterance 'counts as an attempt to get H to do A' he has to recognise it as a request. There is nothing about the utterance that can tell the hearer that the speaker *actually wants* the hearer to do A, because it does not matter for the success of the utterance whether the speaker wants him to do A or not. It is irrelevant for successful communication. But then how does the hearer

[50] Searle's speech act theory depends on there being constitutive rules that differentiate the kinds of things you can do with words, the rule for promising is different to the rule for making assertions, for example. And these rules are what make the speech act the kind of act it is. My view is that all communicative acts share a rule, they all ought to be attempts to get the hearer to recognise one's actual mental state. (The question of what means the speaker employs to achieve her communicative ends - e.g. language, gestures etc. - and how the hearer recognises these, is a related, though different, question.) It does not seem essential to me that speech acts fall into different categories. It is not essential that in order to understand 'What's the time?' one must understand the speech act of questioning as distinct from other kinds of acts. It is only essential that the hearer recognise (all going well) the speaker's desire that he tell her the time.

[51] 1970 pp.66-7.

get in a position to recognise what it is that the utterance counts as? For speech act theory to work, the hearer has to recognise the speaker's illocutionary intention in order for communication to succeed, and yet how this happens, given the lack of necessary connection to the speaker's actual mental state, is a mystery to me.

Searle argues that the hearer can take cues from the context in order to discern what speech act is being performed, and that the speaker could always employ an 'illocutionary force-indicating device'[52] such as saying "I promise" or "I ask you." One could argue, as do Sperber and Wilson, that expectations of *relevance* explain how a hearer is able to recognise a speaker's meaning.[53] However, I believe this kind of answer misses the essential point of my challenge. The claim I am making is that in order for the speech act *to count as* the kind of act that it is, there must be some connection between it and the speaker's mental state. Claiming that there are other ways of coming to recognise the speech act as the kind of act that it is presumes that the act is pre-defined, presumes that it is already imbued with meaning, and the hearer just has to work out what it is. But my question is, how does the act come to have the meaning that it does, how does it come to *be* a request/suggestion/ demand?

It seems to me that the 'counts as' rule, the essential rule, that plays such a key role in Searle's account of speech acts, is *derived from* a more fundamental rule, the rule that utterances ought to be attempts to get the hearer to recognise the speaker's actual mental state. A request is a request because it *ought* to be an attempt to get the hearer to recognise that speaker wants H to do A. And therefore, requests that are not attempts to get the hearer to recognise the speaker's desire that h do A can still *count as* an attempt to get H to do A. If we take out the connection to the speaker's actual desire, then what exactly is it about a 'request' that makes it an attempt to get H to do A? There is nothing about it that says anything about S's actual desire, just as there is nothing about a suggestion or advice that says anything about what the speaker thinks the hearer should do, and there is nothing about a promise that says anything about what a speaker intends to do. Searle cannot make the distinctions he deems so essential without allowing for the *essential normative distinction* between utterances that reflect the speaker's actual attitude, and ones that do not.

---

[52] *Ibid* p.68.
[53] Sperber and Wilson 2005.

Under my account, in contrast, the rule is not a 'counts as' rule, as Searle has it. The rules for communication are rules that can be broken. The hearer is in a position to recognise the speaker's utterance as an attempt to get him to recognise her actual mental state because this is how he ought to understand it. An utterance counts as a request because, if the speaker makes an utterance as an attempt to get the hearer to recognise that she wants him to do A, then his understanding her utterance properly entails that he recognise that she wants him to do A. If she is dissembling, and she doesn't actually want him to do A, or if he takes her to be dissembling, the utterance still 'counts as a request', it can still communicate to the hearer *that* the speaker wants him to do A, or *that* she wants him to think that she wants him to do A, but something has necessarily gone wrong. The hearer has not, upon understanding the utterance, recognised the speaker's actual mental state – he has not recognised her desire that he do A.

Theories of communication that deny a necessary connection between 'sincerity' and the possibility of communication, will struggle, I contend, to maintain that utterances can do the communicative work they are meant to do. Even if the utterance is not meant to communicate the speaker's *actual* mental state, they will struggle to explain how the utterance can come to have any illocutionary force at all. The illocutionary intention is the intention that the hearer take the speaker to be expressing a certain attitude, but if there is nothing about the utterance that links it to the speaker's actual attitude, then it becomes difficult to see how the utterance can *count as* an expression of a particular attitude as opposed to any other.

**

Of course, it is not within the scope of this thesis to argue conclusively for the demise of speech act theory, or all communication theory that denies a necessary connection between successful communication and the speaker's actual mental state. There will be ways of responding to the challenges I have made against them. However, I hope that I have given reasons to doubt that such theories can rest content. Such theories have often been at pains to make sense of the idea that sincere utterances and insincere ones are equally communicative, but this is a path that has a dangerous conclusion. We ought to admit (as I have argued) that insincere utterances are communicative, but that they are not equally correct, or equally technically successful, as sincere ones. In my opinion, the beauty of

constitutive rules has been largely overlooked within speech act theory. There has been a reluctance to define communication in such a way as to disqualify instances of what we clearly call communication (such as when a speaker is deliberately misleading) as communication. The retreat from factivity is felt to be necessary in order not to exclude insincere utterances as being communicative.[54]

Most theories of communication distance themselves from factivity, and yet, I believe, in so doing, they distance themselves from being able to offer a satisfactory account of communication. But it is worth noting that not all theorists have done this. Grice was one of the first to propose a theory of communication. As quoted above, he claimed that : '"A $meant_{NN}$ something by x" is (roughly) equivalent to "A intended the utterance of x to produce some effect in an audience by means of the recognition of this intention".[55] For Grice, successful communication entailed the 'effect' being produced in the hearer. Stephen Schiffer has proposed a theory of meaning which entails a *necessary openness* of intentions on the side of speaker and a 'mutual knowledge' condition, such that both speaker and hearer are aware of the speaker's intentions.[56] I believe his account is more satisfactory as a result of this, but Bach criticises it on the same ground that Searle criticises Grice: of conflating perlocutionary and illocutionary effects, and this has the effect that deceptive cases are not straightforwardly meaningful.[57] The aim of these philosophers is to find the necessary and sufficient conditions for meaning, but none, that I can see, have made use of the idea that a constitutive rule can form the necessary link between the speaker's actual mental state and utterance meaning without the unpalatable consequence of having insincere utterances disqualified as straightforwardly meaningful. (And none have posited that understanding utterances is also a rule governed activity.)

Under my account an utterance succeeds in getting the hearer to recognise the speaker's actual mental state because the rule is that the speaker's utterance ought to be an attempt to get the hearer to recognise her actual mental state, and a hearer ought to, upon understanding her utterance, recognise, thereby, her actual mental state. If the utterance is not an attempt to get the hearer to recognise the speaker's actual mental state

[54] Searle used the notion of constitutive rules, but for him, a constitutive rule was a 'counts as' rule (1970 pp.36-41), and not normative in the sense that it could be broken.
[55] 1957 p.385.
[56] 1972 p.39.
[57] 1970 p.154.

it is a flawed, or incorrect utterance. The speaker has flouted the constitutive rule of communication in that she attempts to get the hearer to recognise her as having a mental state that she does not have. And consequently the hearer fails to obey his rule: he either takes her to have a mental state that she does not have, or rejects her utterance as insincere. Insincere utterances can only succeed in deceiving the hearer because it is a constitutive rule that they be sincere (are attempts to get the hearer to recognise the speaker's actual mental state). That communicative utterances can get a hearer to recognise a speaker's actual mental state is not something special that they can do once communication has occurred (and if certain other conditions obtain), it is what they *do* do when everything goes right. Speakers can cheat and lie, and hearers can doubt a speaker's intentions, but these instances of communication are possible only because they are *necessarily* deviant or flawed instances of communication.

Factiveness has an essential role to play in the possibility of communication: it is *necessary* that a correct or technically successful instance of communication entails that the speaker has the mental state or i/b/d/a that she is attempting to communicate to the hearer, and that the hearer takes the speaker to have the i/b/d/a that she actually does have. If I want some ice-cream, attempt to get you to recognise my desire for ice-cream and you successfully recognise my desire for ice-cream, then communication has succeeded. This only occurs when we have both obeyed our constitutive rules for communication. If anything else occurs, something has gone wrong. Successful communication entails factive understanding, and when this is not achieved, communication has not technically succeeded (though it can still have *occurred*).

The normative distinction is what explains the speaker's communicative intentions (even when her intentions are insincere) and is what explains the hearer's ability to recognise an utterance as an attempt to get him to recognise the speaker's actual mental state. Without this distinction, the possibility of communication occurring at all (correctly or incorrectly) founders. Without this distinction, a hearer is seen to understand an utterance correctly whether or not he sees it as an attempt to get him to recognise the speaker's actual mental state. So even though the hearer is none the wiser about the speaker's actual mental state, he is seen to have been successfully communicated with. Without the normative distinction, it is hard to explain how a speaker communicates her actual mental

state, it is hard to explain her communicative intentions when she makes an utterance, and it is hard to explain how the hearer is able to understand her utterance.

If it is a constitutive rule of communication that an utterance is an attempt to get the hearer to recognise the speaker's actual mental state, then properly understanding it means recognising it as an attempt to get the hearer to recognise the speaker's actual mental state: the speaker's actual mental state is revealed to the hearer. The rule explains how it is that the speaker communicates her mental state to the hearer. When something goes wrong, and the constitutive rules are broken, and the speaker does not communicate her actual mental state to the hearer. But being in a epistemic position to recognise the speaker's actual mental state is not a position that needs to be established with good reasons, it is a position gained merely by being a participant in the institution of communication. Being able to communicate puts one in a position to understand utterances correctly, when nothing has gone wrong, and when one does this, one recognises the speaker's actual mental state.

Communication is not just the expression of attitudes understood as such by the hearer, attitudes that the speaker may or may not have: in making utterances we are not just attempting to give the hearer a reason to think that we have a certain mental state. We, as speakers, are able to reveal our actual mental states to a hearer by making utterances, and, as hearers, we are able to recognise the mental states of a speaker by understanding utterances. We are able to do this because we know how to communicate. This is what happens when communication succeeds under my account. The thin conception favoured by philosophers in speech act theory does not allow this. If the speaker's actual mental state is not linked in any necessary way to the possibility of communication, then it is difficult to see how actual mental states are ever communicated. If the hearer does come to think that the speaker has the i/b/d/a that she actually has, it is not technically *communicated* to him.[58] The epistemic access we have to the mental states of others will always be through a veil of inference, proper understanding of communicative acts does not entail a recognition of the speaker's actual mental state.

On my view, in contrast, the hearer is not presented with the speaker's intention that he take her as having a certain i/b/d/a or with a reason for thinking she has it, he is

[58] As Bach and Harnish say, anything more than the recognition of the speaker's illocutionary intention is 'more than just communication.' 1979 p.16.

presented with her actual i/b/d/a. Properly understanding her utterance entails recognising her as having the i/b/d/a that she is intending for him to recognise her as having, that being her *actual* i/b/d/a. The factivity is entailed in the successful communicative event. When things go well, there is no need for the hearer to have to work from understanding the speaker's utterance to the knowledge that she has a certain mental state: correctly understanding her utterance entails him recognising her as having it. Yes, sometimes things go wrong, sometimes the speaker breaks her rule, and sometimes the hearer breaks his (sometimes without being aware of it). But these situations are *necessarily* incorrect, they are necessarily technically unsuccessful instances of communication. The possibility of them even occurring depends upon the existence of the constitutive rules of communication, and when these rules are conformed to, a speaker communicates her actual i/b/d/a to the hearer.

Communication can occur, and can succeed, because speakers and hearers *know how* to communicate. Knowing how to communicate means being aware of how one ought to behave, one is aware (in the minimal sense) of the rules that guide one's behaviour. Being presented with an utterance, one is not at a loss as to how to understand it: one can understand it as an attempt to get one to recognise the speaker's *actual* mental state. If one wants to communicate something to a hearer, one knows that one ought to communicate one's actual mental state, and that the hearer is in a position to understand it as an attempt to get him to recognise one's actual mental state. Even if these rules can be broken, and even if they can be broken without the participants' knowledge, this does not undermine their existence. Their existence constitutes the very institution of communication.

Conclusion

This chapter constitutes the first part of my explanatory response to the problem of how testimony-based knowledge is possible. The claim that I have aimed to defend is that human communication as an institution depends on the existence of constitutive rules, rules which provide an essential link to factivity. These rules are what make it possible for a speaker to communicate to a hearer her intentions, desires, beliefs or attitudes. The speaker's rule is that she intend to get the hearer to recognise her actual i/b/d/a, and the

rule for the hearer is that he recognise her as trying to get him to recognise her actual i/b/d/a. The rules govern the behaviour of speakers and hearers in a way that constitutes their very ability to communicate: understanding how to communicate is one and the same thing as knowing these constitutive rules. These rules can be broken, and when they are, something necessarily goes wrong: these are instances of deviant or parasitic communication.

What I hope to have achieved in my defence of this conception of communication is that there is a *necessary* link between communication and factivity: the speaker's *actual* mental state features centrally in a satisfactory account of communication. If we understand communication in this way, then, when everything goes well, the hearer, upon understanding an utterance, apprehends directly the speaker's mental state. There is no gap between understanding an utterance and coming to recognise the speaker as having the i/b/d/a that she is trying to get the hearer to recognise her as having through her utterance: properly or correctly understanding her utterance entails recognising her as having the i/b/d/a that she is trying to get him to recognise her as having, and entails that it is her actual i/b/d/a. This factivity establishes that there is no epistemic gap to be crossed by the hearer once he has properly understood a communicative utterance.

In the previous chapter I argued that the majority of approaches to the problem of testimony see their job as one of justification: it is presumed that in order to establish how knowledge through testimony is possible, one must overcome the problem of vulnerability brought about by the indirectness of testimony. It is my contention that this vulnerability problem can (and ought to) be ignored, and that this is because there is no epistemic gap that needs to be crossed by justification. Establishing, as I have attempted to do in this chapter, that successful communication entails the hearer's recognition of the speaker's actual mental state, is part of the argument that will establish that knowledge can be had from an instance of testimony without there being an epistemic gap to be crossed between hearing that p, and coming to know that p. The hearer does not have to work from understanding a communicative utterance to the conclusion that the speaker has the i/b/d/a that she wants him to think that she does.

The possibility of this, of correct communication bringing with it factivity of understanding, is part of the explanation for how knowledge is communicated through testimony. My focus in this chapter has been on the communication of internal mental

states, such as desire, belief, and opinion. But I want to expand this in the following chapter to include the communication of knowledge: an *external* mental state. When one communicates knowledge, as in the case of testimony, one communicates to the hearer one's mental state of knowledge, but more than this, the hearer comes to have that knowledge too. When one communicates an i/b/d/a that one has to the hearer, the hearer recognises the speaker as having the i/b/d/a that she is trying to get him to recognise her as having, but when one communicates knowledge that p, the hearer not only recognises the speaker's knowledge that p, but also comes to have knowledge that p himself. The 'contagiousness' of knowledge is something that requires further explanation.

Asking a question, expressing a desire, or a belief can result in the hearer coming to recognise that desire or that belief in the speaker. Properly communicating one's internal mental state entails that the hearer recognises one's actual mental state. But many communicative utterances (usually in the form of assertions) are attempts at getting the hearer to recognise that some *state of the world* obtains. When I use utterances such as 'The bees are swarming' or 'The capital of Peru is Lima' I am attempting to get the hearer to recognise that the proposition asserted *is true*. This is not, at first glance, an instance of an attempt to get the hearer to recognise my mental state. But this is because knowledge is a different kind of mental state, it is an external, or factive mental state. When one gets the hearer to recognise one's mental state of knowing that p, one can get him to recognise that p. Making sense of this is the content of the next chapter.

Chapter Five: Communicating knowledge

This chapter and the previous chapter together are intended to comprise the third level of Cassam's multi-level approach to how-possible questions. It is at this level that the question of how we get knowledge through testimony can be given a substantive response. It is no longer my aim to remove the obstacle to testimony-based knowledge – this was the aim of the Chapter 3. Once the obstacle (in this case, the obstacle of epistemic vulnerability) has been removed, the question of what it is that *makes* knowledge of this kind possible still remains. In Cassam's terminology, what are the *enabling conditions* that allow for the possibility of testimony-based knowledge? It is my contention that it is the institution of *communication* that allows for the transmission of knowledge, and if we understand how communication works (and what knowledge is), we will be able to understand how it is that a hearer can get knowledge from a speaker, just by understanding her assertion.

In the previous chapter I put forward a particular conception of communication such that a successful instance of communication entails that the hearer recognise the speaker's actual mental state: the speaker, with her utterance, tries to get the hearer to recognise her actual mental state, and the hearer, upon understanding her utterance, recognises her as trying to get him to recognise her actual mental state. This account of communication does not leave an epistemic gap between understanding an utterance and coming to recognise the speaker's mental state. Properly understanding the communicative utterance entails recognising the speaker's actual mental state (the one that she is trying to communicate to him.) What this account of communication establishes is that we can have direct access to the mental states of others via their communicative utterances – we do not have to infer them.

The aim of this chapter is to use this account of communication to explain how it is that communication *of propositional knowledge* occurs, or, in other words, to explain how it is that a speaker can communicate *that p* to a hearer. I aim to explain how testimony-based knowledge is possible. The mental state that the hearer becomes aware of in an instance of testimony is a mental state that is a little different from the speaker's intentions, beliefs, and desires. The mental state that the hearer becomes aware of, recognises the speaker as being in, upon understanding a 'testimonial' communicative act, is the mental state of *knowing*. The distinguishing feature of knowledge as opposed to other mental states is its

factivity (its world-entailing nature); and the distinguishing feature of testimony as opposed to other kinds of communicative acts is that what is communicated is a fact about the world, not an internal state of the speaker. It is the factivity of knowledge that explains how knowledge is communicated through testimony. What is communicated is the *external mental state of knowledge*.

In the first section of this chapter I will outline my positive account. I will argue that if knowledge is the mental state that the speaker is trying to communicate to the hearer, then *successful testimonial communication entails that the hearer come to know what is asserted*. The institution of communication itself explains how it is that testimony is a source of knowledge for a hearer. The hearer coming to know what is asserted is not something that requires provision of epistemic reasons, or justification: the hearer has, when all goes well, direct access to the speaker's knowledge that p, and thus direct access to the fact that p. Successful communication fully explains how it is that a hearer can get knowledge from a speaker. In the second section of this chapter I will attempt to respond to what I believe will be the most powerful objections to my account.

Section One: How to communicate knowledge

The types of communicative utterances that I focused on in the previous chapter were utterances that were attempts to get the hearer to recognise the speaker's internal mental state. The woman who asks the man to close the window uses her utterance to get the man to recognise her *desire* that he close the window. The opera lover who attempts to get the hearer to recognise her love for opera uses her utterance to get him to recognise a certain *attitude* or opinion that she has. When testimony occurs, however, what the speaker is trying to communicate to the hearer is not (necessarily) her intention/belief/desire/attitude, but *that a certain fact obtains*. She is not trying to get the hearer to recognise her *internal* mental state, she is trying to get him to recognise that a certain proposition is true. The intention behind an instance of testimony is to share or communicate propositional knowledge.

In an instance of testimony, it is not the speaker's intention to communicate to the hearer anything about her internal mental state, it is her intention to communicate to the hearer *that p*. Unlike instances of communication where the speaker wishes to

communicate her desire or her intentions, the hearer recognising her *internal* mental state is not her intention. She intends for the hearer to recognise the truth of what she asserts, and thereby come to be in the same mental state as her: she intends for him to come to know that p, as she does, upon understanding her utterance. This distinguishing feature of testimony is explained by the fact that what makes knowledge different from other mental states is its *factiveness* or its *externality*.[1] Knowing that p entails that p is true, and it is the speaker's intention that the hearer recognise *that p is true* upon understanding her utterance properly. Her correct assertion that p entails that p is true, and therefore a correct understanding of her assertion that p entails recognising that p is true. It is in this way that *the hearer can come to know that p* based merely on understanding the speaker's communicative utterance: her assertion that p.

The type of mental state that the speaker is trying to communicate to the hearer in the case of testimony explains the fact that the speaker's intention is not to get the hearer to recognise her *internal* state (e.g. i/b/d/a), but to recognise a certain state of affairs as obtaining. It also explains how the hearer can come to know that p just by hearing and understanding an assertion that p. The mental state is knowledge, and knowledge is a factive mental state, an external mental state: knowing that p entails that p is true. The factivity of such a state (the truth of what is asserted) is what the speaker is trying to get the hearer to recognise: when the hearer properly understands an assertion that p, he is recognising that p is true because the speaker's mental state of knowing that p entails that p, and what he is recognising, when he understands the utterance, is the speaker's actual mental state. The recognising of the mental state of knowing is the recognising of p as true. The hearer's coming to know that p is what the speaker intends, and this is the intention that is fulfilled when the hearer recognises it.

One thing that distinguishes testimony from the kind of communicative utterances that I was discussing in the previous chapter, is that successful communication does not *seem* to require that the hearer recognise the speaker as having a certain mental state. The communicative intention behind my utterance of 'The capital of Peru is Lima' does not seem to entail an intention that the hearer come to recognise anything *about me*. It seems more accurate to say that my communicative intention is that the hearer come to recognise *that*

[1] 'Factiveness' and 'externality' are meant to point to the same feature of knowledge: it's world-entailing nature.

the capital of Peru is Lima. My intention, unlike with 'non-testimonial' communication, is to get the hearer to recognise that some state of affairs obtains. It is my intention, when I assert that p, to get the hearer to recognise *that p*. It is my intention, in other words, that the hearer, upon understanding my utterance, can come to recognise that p is true. But, although testimony does not involve the intention that the hearer recognise something *about the speaker*, he is still recognising *her mental state*. It is just that the mental state is an external, or factive, or world entailing mental state, and therefore recognising it does not (necessarily) entail recognising something about the speaker.

It should be noted that states of the world communicated through testimony can include facts about the speaker. Assertions such as 'I am hungry', or 'I have a headache' or 'I really enjoyed that film' are still testimonial utterances. They are instances of telling the hearer *that* I am hungry, or *that* I enjoyed that film. (The hearer recognises that I am hungry in the same way that he recognises that Lima is the capital of Peru.) What this means is that one can use testimony to convey one's i/b/d/a, but it differs from non-testimonial communication in that the speaker is trying to get the hearer to recognise *the truth of the proposition* that she would like ice-cream (for example), and not to recognise *her desire* for ice-cream. (Though they may, in practice, amount to the same thing – saying 'I would like some ice-cream' is often just another way of saying 'Please may I have some ice-cream?')

The distinctive communicative intention (to communicate something about the world, and not the internal state of the speaker) does not show that testimony is fundamentally different to the kinds of communication that I discussed in the previous chapter. Testimony is a form of communication like any other. What makes it different is that the mental state that the speaker is trying to get the hearer to recognise is the mental state of *knowing*. The mental state of knowledge differs from other mental states in its *factiveness*. What this means is that it is an 'external' mental state, not internal in the way that beliefs, desires and intentions are. It is the distinguishing feature of knowledge as a mental state (that it is factive) that can explain the apparent distinctiveness of the speaker's communicative intention. The fact that the speaker's intention behind the 'testimonial' utterance (usually assertion) is to get the hearer to recognise that p, can be explained by the fact that what she is trying to get the hearer to recognise is her mental state that *entails that p*.

My intention behind the (testimonial) utterance 'The capital of Peru is Lima' is to get the hearer to recognise that the capital of Peru is Lima, but to do this he must recognise *my*

mental state of knowing that the capital of Peru is Lima. What he recognises, upon understanding my assertion, is the factivity of my mental state. He comes to know that the capital of Peru is Lima via recognition of my intention to get him to recognise this *in* his recognition of my (external) mental state of knowing. What he is recognising when he recognises my mental state *is* that the capital of Peru is Lima. It is still a case of the speaker getting the hearer to recognise her mental state, and yet the distinctive nature of the mental state is that it is factive, and it is its factivity that the speaker is trying to get the hearer to recognise. The *truth* of what is asserted is what the speaker is trying to communicate to the hearer. As with the communication of non-testimonial utterances, the hearer's understanding of the assertion does not leave it open whether or not the speaker is in the mental state she is trying to get him to recognise her as being in. The act of understanding her utterance entails a recognition of what she asserts as true: it is one and the same cognitive event. Correctly understanding an assertion *is* recognising it as true.

Testimony is still an instance of communication in the way that I have described in the previous chapter. The speaker uses her utterance in an attempt to get the hearer to recognise her (as being in a) certain mental state, and the hearer understands her utterance when he recognises her (as being in that) mental state. However, when that mental state is knowledge of p, it is p's *being true* that is relevant to the speaker, and to the hearer. It is not the speaker's *internal* mental state that she is trying to communicate to him: it is an external mental state, a state that encompasses the fact that p. And it is not the speaker's internal mental state that the hearer comes to recognise upon understanding her utterance, but what her external mental state entails – that p is true. A speaker communicates *that p* through the distinctiveness of the mental state that she is communicating.

The constitutive rules for communication, as I set them out in the previous chapter, are:

> **Speaker's rule**: The speaker's utterance must be an explicit attempt to get the hearer to recognise her (as having an) intention/desire/belief/attitude (or relevant mental state), the i/b/d/a that she actually does have;
>
> **Hearer's rule**: The hearer must recognise, upon hearing the speaker's utterance, that it is an explicit attempt to get him to recognise her (as having an) intention/desire/belief/ attitude, the i/d/b/a that she *actually does have*. (This rule

entails factive understanding: the state he takes her to be in must be the state that she is trying to get him to recognise her as being in AND she must actually be in it.)

When the relevant mental state is knowledge, the picture *appears* a little more complex:

> **Speaker's rule**: The speaker's utterance must be an explicit attempt to get the hearer to recognise *that p* in his recognition of her knowledge that p, knowledge that she actually does have;[2]
>
> **Hearer's rule**: The hearer must recognise, upon hearing the speaker's utterance, that it is an explicit attempt to get him to recognise that p *in his recognition of her knowledge that p*, knowledge that she actually does have. (This rule entails factive understanding: when he recognises her knowledge that p, she actually does know it: and thus so does he.)

The successful instance of testimony entails that the speaker knows what she asserts,[3] and that the hearer recognises this knowledge, and thus also come to know it. A successful instance of testimony entails that the knowledge that the speaker intends to share is shared: the hearer comes to know what is asserted.

The picture appears more complex because it seems to suggest that the hearer must recognise two things: the speaker's knowledge that p and that p. However, these are not two different things. The speaker's knowledge that p is her cognitive grasp on the fact that p. What makes it knowledge (as opposed to any other mental state) is its world entailing nature. The hearer, upon recognising her knowledge, does not need to recognise anything more than the fact that p obtains. The mental state that he is recognising is not her belief that p, or her good reasons for believing it. Knowledge is not a state that can be reduced into conceptually distinct components (as internalists holds). The speaker's knowledge is an

[2] What I mean by 'recognise that p in his recognition of her knowledge that p' should hopefully become clearer as the chapter progresses. The hearer's recognition of p and his recognition of the speaker's knowing that p are not separate cognitive occurrences: to understand the utterance 'The capital of Peru is Lima' *is* to come to know that the capital of Peru is Lima.

[3] One may be familiar with Timothy Williamson's knowledge account of assertion (2000, Chapter 11). He argues that there is a constitutive rule for assertion, and it is that 'one must: assert p only if one knows p' (p.243). His account, I would say, complements my views, though his focus is on assertions in general, and not on communicative utterances in particular – which have a specific intention. Charles Pelling (2013) argues that the norm for telling is not derived from the norm for assertion; but perhaps, and this would be a subject for another paper, the norm for assertion is derived from the norm for telling.

irreducible state, it is a state that entails the truth of p, and in recognising this state, what the hearer is recognising, when all goes well, is the truth of p.

Testimony, as a kind of communicative event, does not entail any more rules than the two I have already stipulated. It is not that, in the case of testimony, the speaker is governed by a rule other than the rule that she ought to attempt to get the hearer to recognise her as being in a state she actually is in. And understanding a testimonial utterance does not require any more of the hearer than that he recognise the speaker as being in the state that she is trying to get him to recognise her as being in, and that this is her actual state. The original constitutive rules are not changed. What is changed is the nature of the mental state that is being communicated: its distinctive nature of factivity is what explains the shift in speaker intention. The speaker wants to communicate *that p* to the hearer, rather than communicate her internal mental state, but it is still her mental state of knowing that p that is being communicated.

It is an essential point that the speaker is not communicating to the hearer *that she knows* that p with her assertion that p, any more than a speaker is communicating *that she desires* to know whether p when she asks whether p. In the latter case the speaker is trying to get the hearer to recognise *her desire* to find out whether p, not to recognise *that* she desires it. And likewise, in the former case the speaker is trying to get the hearer to recognise *her knowledge* that p, not *that* she knows that p. The proposition that she wants him to come to recognise as true is *that p*, not that she knows that p. Properly understanding a non-testimonial utterance entails recognising the speaker's actual i/b/d/a (not the proposition *that* she has a certain i/b/d/a). Properly understanding a testimonial utterance entails recognising the speaker's actual mental state of knowledge, without the hearer having to recognise *that* the speaker knows what she asserts. The proposition that the hearer recognises as true is the proposition that the speaker is trying to get the hearer to recognise is true: the hearer recognises, upon (properly) understanding an assertion 'p', that p.

It should be noted that the rule does not entail a requirement on the *form* or *content* of the utterance itself. One can testimonially communicate that p to a hearer without uttering an assertion with the content 'p'. In order for the communication to be a success, the speaker must be trying to get the hearer to recognise that p in his recognition of the speaker's knowledge that p, and this does not necessarily entail that the form of the

utterance is the proposition 'p'. Sarcastic utterances can constitute testimonial communicative success. It is also possible for people devise a code in order to communicate secrets in public, whereby the form or content of the utterance bears no relationship to what is being communicated.[4] And one can make mistakes such that the literal meaning of the words bears little resemblance to the state of affairs that the speaker is trying to communicate. But as long as the hearer is still able to tell what the speaker is trying to communicate, then success testimonial communication can occur.[5] Successful communication requires only that the hearer recognises what the speaker intends for him to recognise in his understanding of her utterance.

Perhaps the best way to clarify this view of testimony is through an example of an everyday instance of testimony. I am at the library and I ask the librarian what time the library closes. She tells me, 'The library closes at ten.' Under my account, the librarian obeys the constitutive rule of communication only if she knows that the library closes at ten. Her utterance is an attempt to get me to recognise that the library closes at ten in my recognition of her knowledge that it closes at ten. I understand her response, I recognise her utterance as an attempt to get me to recognise that the library closes at ten in my recognition of her knowledge that it closes at ten, and I, thereby, do recognise that the library closes at ten. My understanding of her communicative utterance entails a recognition of the *truth* of what she asserts. If I am still unsure of what time the library closes, or if I come to believe it closes at ten, but it actually closes at nine, then I have not successfully understood her utterance. Understanding her utterance properly (in the sense required for successful communication) requires that she actually does know what she asserts: understanding a 'testimonial' utterance (assertion) properly entails *coming to know* what is asserted.[6]

The factiveness (or world entailing nature) of knowledge is what distinguishes it as a mental state. It is also what distinguishes testimony from other kinds of communication. It

[4] As mentioned in footnotes in the previous chapter, there must also be a common ground recognised such that the speaker can expect her utterance to be the means to communicate the mental state that she intends it to. Language is the most common means, and language has constitutive rules of its own that allow it to function as a means to communication.

[5] A consequence of this is that an assertion that p (under my account) is a testimonial communicative utterance used with the intention of communicating p, and not an utterance that takes the form 'p'.

[6] The intuition that proper understanding does not require acceptance of an assertion as true is a strong one. But I will address this issue in Section 2 (Objection v.).

is the world-entailing nature of knowledge that is communicated with testimonial utterances. Another way of putting this is that it is the *truth* of what is asserted that is communicated through testimony. The hearer does not need to form any beliefs about the speaker's *internal* state at all. For example, if a person were to shout 'Fire!', what is being communicated is that there is a fire, and not anything about the speaker, the hearer might not even know who said it. However, the speaker *has to know* that there is a fire in order for this to be communicated to the hearer. And in order for the hearer to properly understand the utterance 'Fire!' (and thus come to know that there is a fire) he must recognise that there is a fire in recognising the speaker's knowledge that there is a fire. It is the speaker's mental state of knowing that there is a fire that guarantees that there is a fire: her mental state of knowing entails that there is a fire.

The factiveness of knowing explains why, in a successful instance of testimony, the hearer cannot simply recognise the speaker's knowledge, he must come to have knowledge too. I can recognise a speaker as believing that p, or having the opinion that p, or having a certain desire or a certain intention, but this will not bring about in me the same belief, opinion or desire. However, when I recognise a speaker's knowledge that p, I come to know it too. This is because knowledge is factive (world entailing), and recognising a speaker's knowledge that p means recognising that p is true, and thus coming to know it too. This is not to say that knowledge is *inferred* from recognition of the speaker's knowing that p (as I have already said, the hearer may not form any beliefs *about the speaker* at all). The hearer recognises that p is true upon understanding an assertion (testimonial utterance), [7] because if the utterance is correct (the speaker is obeying her constitutive rule), p *is* true. (If the utterance is correct, the speaker has knowledge that p.)

It is worth mentioning here that the idea that assertions are utterances governed by the 'knowledge rule' is an idea espoused by many, especially since Williamson defended the view in Chapter 11 of his book *Knowledge and its Limits*.[8] However, just as with

[7] I will use 'assertion' interchangeably with 'testimonial communicative utterance' from now on, even though not all testimonial communicative utterances will look like (what we would normally call) assertions (e.g. a nod), and not all (what we would normally call) assertions are testimonial communicative utterances, (e.g. 'It's a lovely day!' to someone standing right next to you.) Furthermore, under my account, it is possible for one to assert that p without uttering the proposition p. For my purposes, an assertion that p will be a testimonial communicative act performed as an attempt to communicate that p.

[8] 2000. For the knowledge rule, see Turri (2011), Ball (2014), Benton (2011, 2012), DeRose (2002), Kelp (2013), Reynolds (2002); for other epistemic rules (e.g., a 'justified belief rule') see Kvanvig (2009), Douven (2006), Lackey (2007), (or the 'truth rule') Weiner (2007).

communicating internal mental states, the speaker's rule on its own would not be enough to explain how knowledge can be transmitted through testimony. Sanford Goldberg explores the possibility of this in his paper 'The Knowledge Account of Assertion and the Nature of Testimonial Knowledge',[9] and finds it wanting. This is because, I contend, in order for the knowledge rule of assertion to explain how the hearer gets testimony from the speaker, the hearer has to *take the speaker as conforming to the knowledge rule*. The hearer must be able to 'reliably distinguish cases of assertion that conform to [the knowledge rule] from cases of assertion that do not.'[10] If, however, the interpretation of utterances is guided by the rule that one ought to recognise the speaker's testimonial communicative utterance that p as an attempt to get the hearer to recognise her actual knowledge that p, then the hearer (just by being a participant in the institution of communication) is in a position to understand an assertion as conforming to the knowledge rule, unless there are reasons to think something has gone wrong.

Breaking the rules

The speaker can break her rule by asserting that p when she does not know that p. She can do this intentionally, by lying, or she can do it unintentionally, by being mistaken. When a speaker is lying she is trying to get the hearer to believe that p, even though she herself does not believe that p. The speaker's intention is to get the hearer to think/believe that something is true, when she (the speaker) knows that it is not. For example, a woman lies to her husband about her whereabouts. She mendaciously says to him, 'I was at work all evening'. This is an instance of the speaker breaking the constitutive rule of communication. Her communicative intention can be fulfilled if her husband comes to believe her: he takes her to be *obeying* her rule. However, communication does not technically succeed in this instance, even if the husband does believe her: he believes that she was at work all evening because he has taken her to be in the mental state of knowing that she was at work all evening (he believes he is recognising her knowledge that she was at work), but she was not in that state. His believing her occurs only because he mistakenly

[9] 2009, in Greenough and Pritchard (eds.).
[10] *Ibid.* p.62.

believes her utterance to be an attempt to get him to recognise that she was at work in his recognition of her (*actual*) knowledge that she was at work. He cannot get knowledge from her in this instance, this is a necessarily deviant or parasitic instance of communication.

A speaker can also break her rule by asserting something that she thinks that she knows, but she happens to be wrong about it. Say, for example, a mother says to her son 'There are nine planets in our solar system.' However, scientists have recently agreed that there are only eight planets in the solar system, and thus her assertion breaks the knowledge rule. False assertions that arise from the speaker's sincere belief that p, and successfully get the hearer to believe that p, are *not* instances of successful communication. This may seem counterintuitive to some, as the speaker believes that p, makes an assertion as an attempt to get the hearer to believe that p, and succeeds in getting the hearer to believe that p. However, as with the case of lying, the speaker can only succeed in getting the hearer to believe what she asserts via his recognition of her intention to communicate her *knowledge* to him. The hearer does not get knowledge from the speaker in this instance, something has gone wrong.

The hearer can break his rule by either rejecting the utterance as an attempt to get him to recognise the speaker's knowledge that p (he doubts that she has knowledge that p – he doubts that she is obeying her rule) or he can take it as an attempt to get him to recognise the speaker's knowledge, but he is wrong, and the speaker does not know what she asserts. In the first instance he may suspect that she is lying, or have reason to think that what she asserts is false. (Perhaps the son has recently learnt at school that there are only eight planets in the solar system.) Defeaters,[11] or reasons to believe that something has gone wrong, can bring the hearer to break his rule. He can reject the speaker's utterance as an attempt to get him to recognise her as being in a mental state that she actually is in, that is, he can think that she is not actually in it. In this instance too, something has gone wrong, he has not properly understood her utterance. Properly understanding it in the way required

[11] Elizabeth Fricker distinguishes between two kinds of defeating conditions: a defeating condition may 'defeat a proposition, in the sense that it constitutes strong evidence for the falsity of that proposition'; or it may 'cancel the right to presume that proposition to be true – being a circumstance which indicates that the proposition may not be, or cannot be assumed to be true, rather than being definite evidence for its falsity' (1994 p.142). It is something like this that I have in mind that can bring the hearer to reject an assertion as a speaker's attempt to get him to recognise her knowledge.

for successful communication entails that he recognise her as being in the state that she is attempting to get him to recognise her as being in, and that it is her actual mental state. When the son does not take his mother as knowing there are nine planets in the solar system, something has gone wrong.

He can also break his rule by taking her to be trying to get him to recognise her as being in the state of knowing what she asserts, and yet, unbeknownst to him, she does not know what she asserts. This can happen when the speaker successfully lies to him: she gets him to think that something is true that she knows is not true. The husband who believes his lying wife has broken his rule. It can also happen when the hearer takes her to know something that she also thinks that she knows. For example, the son who believes his mother when she says that there are nine planets in our solar system has broken the hearer's rule. Another example could be the librarian who tells me the library closes at ten, but she's forgotten that it closes at nine on a Saturday, and that day happens to be a Saturday. I come to believe that the library closes at ten that day, based on my (mistaken) 'recognition' of the librarian's knowledge that it closes at ten. But she does not know it, and therefore neither can I. Something has gone wrong: the speaker has broken her rule, and thus the hearer cannot obey his.[12]

**

The constitutive rules for communication are 'factivity entailing'. When successful communication occurs the hearer comes to recognise the speaker's actual mental state just by understanding her utterance. In the case of testimony, that mental state is knowledge, and thus the hearer's understanding of the communicative utterance entails recognition of what is asserted as true: and thus the speaker's intention - to communicate what is asserted (to get the hearer to recognise that what she asserts is true) - can be effected just by the hearer's recognition of her communicative intention. The fact that the speaker's knowing that p is factive allows for her (correct) assertion to be a means to communicating that p to the hearer simply via his recognition of her mental state (his correct understanding of her utterance).

[12] As I pointed out in the previous chapter: if the speaker breaks her rule, the hearer necessarily breaks his. His rule entails a recognition of the speaker's intention to get him to recognise her actual mental state, and if that is not her intention, then he cannot correctly recognise it as such.

The factivity of knowing adds another layer to the picture I painted in the previous chapter. The *truth* of the proposition asserted in an instance of testimony is a part of the act of properly asserting, and understanding it as true is part of the act of properly understanding an assertion. This means that if what is asserted happens to be false, then something *necessarily* goes wrong with that instance of testimony, even if the participants in the event are oblivious to this. The factiveness of knowledge places constraints on what constitutes successful assertoric (testimonial) communication, even if the transgression of these constraints is sometimes not discernible to speaker or hearer. A speaker may be wrong about what she thinks she knows, and a hearer may, as a result, come to believe something that he does not know. It is possible for all concerned to be under the illusion that successful communication is taking place, and yet it not be.

But this is only to be expected if what is being communicated are facts or states of affairs. People can be wrong about what they think they know, and thus they can believe that they are successfully communicating that p, but not be. A few years ago a teacher would (blamelessly) teach her students that there were nine planets in the solar system. Scientists would not have disagreed with her. And yet, further studies revealed that Pluto was not technically a planet, and thus we have since discovered that there are only eight planets in the solar system.[13] The teacher's communication with her students, her telling them something false, was necessarily flawed. Even if it had not been discovered that Pluto was not a planet, it would still be flawed. Successfully communicating that p, like seeing that p, or remembering that p, is factive. One can only do it if p obtains.

The idea of constitutive rules must allow for there to be a gap between obeying the rules and knowing that one is obeying the rules. It does not even need to be possible for one to discover that one is breaking the rules, but this does not change the fact that the rules govern nevertheless. In order to be able to communicate, speakers and hearers are necessarily governed by the constitutive rules of communication. We know the rules, even if we cannot articulate them. Knowing how to make a communicative utterance, and knowing how to understand one, entails knowing that a testimonial utterance (for example) ought to be an attempt to get the hearer to recognise that p is true only if one knows that p is true,

[13] This is perhaps a naïve view of the situation, it might be that the definition of 'planet' became more precise, and so there *were* nine planets and now, given the new definition, there are only eight, but I will leave this possibility aside for now.

and a testimonial utterance ought to be understood as true.[14] And not taking it as such means that something has gone wrong.

Section Two: Some objections and replies

In this section I consider what I believe will be the most powerful objections to my account. The first part is dedicated to dealing with how other theories of communication make sense of assertion as a kind of speech act. I will argue, in parallel to the previous chapter, that such accounts are too thin. If the hearer coming to believe the assertion is left out of the communicative act, then we are left with an epistemic gap too wide to be crossed. The second part of this section will be devoted to responding to objections to my account that are not to do with theory of communication directly, problems such as the intuition that my account is a promotion of gullibility.

One of the most popular alternatives to this account of communication would be the one offered by speech act theorists. Speech act theorists tend to group speech acts into types, and 'assertives' or 'assertions' are usually considered to be one of the most important types of speech act. For John Searle, assertions are identified by their essential rule: assertions 'count as an undertaking to the effect that p represents an actual state of affairs.'[15] (The sincerity rule is that S believes that p.[16]) For Bach and Harnish: 'S asserts that p if S expresses: i. the belief that p, and ii. the intention that H believe that p.'[17] The main point of contention between us will be the fact that, for me, a successful instance of testimonial communication entails that the hearer get knowledge from the speaker. For the speech act theorists, a successful assertion, correctly understood by the hearer, need not even be believed by the speaker or the hearer.

Searle claims: 'In the case of illocutionary acts we succeed in doing what we are trying to do by getting our audience to recognize what we are trying to do. But the 'effect' on the hearer is not a belief or response, it consists simply in the hearer understanding the

[14] Knowledge of the rules, as I have said before, is not *awareness* of the rules. It consists in the ability to participate in the institution of communication.
[15] 1970 p.66.
[16] *Ibid.*
[17] 1980 p.42.

utterance of the speaker.'[18] For Searle, and for almost all other communication theorists that I have read,[19] the hearer actually believing what the speaker asserts is considered to be a 'perlocutionary'[20] effect of the communicative act. The perlocutionary intention is the intention that the speaker has beyond the basic communicative intention. For example, the (non-assertoric) utterance 'Please close the door' is properly understood, and thus the illocutionary intention fulfilled, when the hearer recognises what the speaker means, (that is, what act the speaker is performing with her words). The perlocutionary effect is the hearer actually closing the door, the speaker's perlocutionary intention is what she intends to happen once the communicative act is completed. The 'perlocutionary' intention and effect is 'extra-communicative.'

The idea is that the equivalent intention with assertions is that once the hearer has understood a speaker's assertion, he may come to believe it. In the same way that the utterance, 'Please close the door' may bring the hearer to close the door, the utterance 'Lima is the capital of Peru' may bring about in the hearer the belief that Lima is the capital of Peru. The hearer's resultant belief in what is asserted is not something that needs to be secured in order for successful communication to occur. For Searle, the hearer must recognise that the speaker's assertion is 'an undertaking to the effect that p represents an actual state of affairs.'[21] Once the hearer recognises this, communication has been successfully achieved.

For Bach and Harnish: 'In the case of statements the speaker expresses two attitudes: belief in a certain proposition and the intention that the hearer believe it as well. That is to say, for S's utterance of e to be a statement that P, S must R-intend [reflexively intend] H to take the utterance as reason to think (a) that S believes that P and (b) that S intends for H to believe that O. Correlatively, for H to understand that S is stating that P in uttering e, H must take S's utterance of e as R-intended to be reason to think (a) and (b). For a statement

---

[18] 1970 p.47.

[19] Even H.H. Clark, whose particular focus is the importance of the hearer's role in communication, claims 'Perlocutions aren't part of understanding itself. … [My son] could have understood my assertion about the book, but not believed me.' (1996 p.133) McDowell would be an exception, though I am not sure that he would count as a communication theorist. He does not talk specifically about the illocutionary and perlocutionary effect of assertions (or any other speech act), but I presume he would disagree with the standard model that acceptance of an assertion as true is a perlocutionary effect. See 'Meaning, Communication and Knowledge' 1980.

[20] Originally J.L. Austin's terminology, 1962.

[21] p.66.

to have been made and to *be successful as an act of communication*, it is not necessary that H actually think that S believes that p or that H believe that P himself. These would be perlocutionary effects of S's utterance.'[22] The hearer's understanding consists simply in his recognition of the kind of act that is being performed, he does not need to come to have the belief proffered, or even to recognise the utterance as sincere.

For me, these accounts are hopelessly thin. They are too thin to explain the speaker's intention, and they are too thin to explain how it is that testimonial communication can succeed, when it succeeds in the intuitive sense of *communicating to the hearer that p*. Bach and Harnish feel compelled to claim that assertions have a *dual* intention. The speaker not only wants the hearer to recognise that she believes that p, but she also wants him to recognise that she wants him to believe it too. My first question is, how could the speaker's intention to get the hearer to believe that p ever be fulfilled? It is certainly not fulfilled just by the hearer recognising the intention, and that is if he takes the speaker to actually have the belief and the intention, which he *need not do* in order for communication to succeed. So the hearer has to get from recognising that the speaker is 'expressing' a certain belief and a certain intention (though how he is meant to recognise that she wants him to think this is inexplicable), to believing or accepting that she actually does have that belief and that intention, to fulfilling that intention and coming to believe it too. How does this ever happen? And how can it be the speaker's intention that such a thing happen every time she makes an assertion? And the possibility that such a belief (the hearer's resultant belief that p) could count as *knowledge* seems (at least at first glance) extremely slim.

Firstly, why is the speaker's expressing of a certain belief a reason to think that she has it? Expressing the belief does not entail that she has it, but is meant to *count as a reason* to think that she has it. But there are plenty of situations in which a speaker may want a hearer to think that she believes something when she does not believe it. And secondly, even if he does think that she believes it, and does see that she wants him to believe it too, (and this is a big 'if') how does this count as a reason for *him* to believe it? Of course, it is not impossible, seeing that someone believes that p may turn out to be a good reason to believe it too, but there would need to be a lot of philosophy done to explain why this is the case. And even if it turns out that seeing that a speaker believes that p *is* a good reason to

[22] p.16 (emphasis added).

believe it, is this really enough to explain how *knowledge* can be had?[23] Explaining how the hearer comes to *know* that p based on hearing a speaker assert it, when hearing it asserted leaves it open whether or not the speaker even believes it (never mind justifiably believes it, or knows it), is not easy.

For Bach and Harnish, it is not their problem: 'anything more [than recognising the speech act for what it is] is more than just communication.'[24] But, as I argued in Chapter 4, a satisfactory account of communication ought to explain how it is that one communicates to a hearer one's mental state, and, even better, it would explain how one communicates that p through an assertion 'p'. Communicating knowledge is one of the most fundamental functions of language, and a communication theory that fails to account for it, fails (in my book) to be satisfactory. But not only this (even if explaining how communicating that p occurs is not their problem), important parts of their own account are left unexplained (and, under their jurisdiction, inexplicable): an assertion is necessarily, by their own definition, an expression of an intention to share a certain belief, and yet there is no suggestion of how that intention is meant to be fulfilled. The illocutionary intention of an assertion is (amongst other things) the intention that the hearer recognise the speaker's intention to share her belief, and yet how this belief is meant to be shared goes unexplained.

Under my account, in contrast, the communicative intention is fulfilled when nothing goes wrong. If both speaker and hearer conform to their constitutive rule (which they know how to do because they know how to communicate), the speaker not only shares her belief that p, but shares her knowledge that p. Something has gone wrong when the hearer fails to come to know what is asserted. If the hearer believes what is asserted, and yet does not know it (it happens to be false), something has gone wrong. If the hearer sees that the speaker believes what she asserts, but not that she knows it (he thinks she is mistaken), something has gone wrong. And if the hearer doesn't believe that the speaker believes it (he thinks she is lying), something has gone wrong. These may still count as instances of communication, but they are necessarily flawed, necessarily incorrect or unsuccessful. When everything goes well, a successful instance of testimony entails that the hearer comes to know what is asserted: the speaker obeys her rule, and asserts what she knows, and the

[23] This is essentially the problem I raised with justificatory responses to the problem of testimony in Chapter 2.
[24] 1980 p.16.

hearer obeys his rule and recognises her as trying to get him to recognise that what she asserts is true via his recognition of her actual knowledge of what she asserts.

The speech act theorists maintain that belief, and not knowledge, is the relevant mental state in defining assertions. An assertion that p is considered to be an *expression* of the speaker's belief in p (as well as some other intentions perhaps), even if the speaker does not believe that p.[25] This may appear more palatable than my account in the first instance because there may be worries about knowledge being a mental state, *per se*, or worries that, from the speaker's perspective, there is nothing to distinguish belief from knowledge.[26] However, it is only knowledge, and not belief, that has the required 'communicability.' Michael Welbourne puts it this way:

> 'PC (Principle of Communicability): If a hearer (H) believes that a speaker (S) knows that P, then H believes that H knows that P.'[27]

It is only knowledge that has this feature. A hearer believing that a speaker believes that p will not result in the hearer believing that p, but one cannot think that a speaker knows that p without coming to believe it oneself.[28] This is because believing that the speaker knows that p entails believing that p is true.[29]

This is not to say that the hearer gets knowledge just by thinking the speaker has knowledge. The speaker has to actually have it. But if the speaker does have knowledge, then the hearer can have it too, just by recognising the (external) mental state of the speaker. This special feature of knowledge is a necessary component of the explanation for

---

[25] *Ibid.* p.15, and Searle 1979 p.65.
[26] I deal with these concerns later on.
[27] 2001 p.110.
[28] Robert Brandom's account of assertion involves an in depth explanation of this feature of assertions from within the internalist perspective. 'In taking someone to be a knower, one *attributes* a *commitment, attributes entitlement* to that commitment, and *acknowledges commitment* to the same content oneself' (p.202). This corresponds to taking the speaker to *believing* the content of the assertion, being *justified* in believing the content, and the content being *true*. In other words, taking the speaker to have knowledge is taking the speaker to have justified true belief. Though it is not clear why it is necessary to take the speaker as being *justified* if one is already entitled to take the assertion as true – unless the being justified is meant to establish that he (the hearer) can take it as true - in which case, my arguments in Chapter 2, Section 2 apply. It is also a challenge for any internalist to explain how it is that a speaker's assertion can communicate the *truth* of the proposition uttered, seeing as, for them, it is necessarily the case that the *truth* of what one knows is simply a state of affairs wholly distinct from the internal state of the speaker.
[29] Welbourne's thesis is that the *essential communicability* of knowledge ought to be the basis of how we understand the concept of knowledge, and that we ought to reject 'belief-theoretical' accounts of knowledge (what I would call internalist accounts). Very influential for him, and for me, is a paper entitled 'Telling the Facts' by Zeno Vendler (1980). In it Vendler argues that the concept of knowledge (and *not* justified true belief) can explain how assertions can 'carry the factive burden' (p.283).

how it is that knowledge can be communicated through an assertion. Thin accounts of assertion, ones that focus on belief instead of knowledge, cannot satisfactorily account for the hearer's taking on of the proffered belief. For the speech act theorists, the hearer recognising the speaker as believing that p has to somehow count as a *reason* for the hearer to take on that belief too. But why should this be so? The fact that the speaker has a certain belief might (but might not) speak in favour of that belief, but is this really enough to take on the belief oneself? And how does it help if this belief is combined with an intention that the hearer take on the belief? This does not (*prima facie*) make taking on the belief a reasonable or epistemologically warranted action. However, if I take the speaker to have *knowledge* of something, I cannot at once believe this, and believe that what the speaker knows is false. If the speaker knows it, then it is true: taking the speaker as knowing that p *entails* taking p to be true.

The problem with this for the speech act theorists might be the same as it would be for epistemologists of testimony: on what grounds can the hearer presume that the speaker has knowledge? If the speaker does have knowledge, then, yes, what she believes is true. But how can the hearer know that she has knowledge? Her assertion will look exactly the same if she thinks she has knowledge, and yet does not. The truth of her belief (what distinguishes belief from knowledge) is not something that is discernible to the hearer upon recognising or understanding the speaker's assertion. If it is sometimes not even discernible to the *speaker* whether or not she has knowledge, how can we expect the hearer to simply recognise that she has knowledge upon understanding an assertion? Surely, *at best*, all that the hearer could recognise is that the speaker *thinks that* she knows what she asserts. And this is the same as recognising that she believes it (and perhaps is justified in believing it[30]) but no more than this.

The idea that the hearer can just know that p upon hearing and understanding an assertion that p might sound preposterous to those committed to an idea of knowledge that

[30] Would the belief being 'justified' do enough to make a significant difference to the explanatory satisfactoriness of speech act theory? After all, if a speaker thinks that she knows something, the thought goes, she must have good reason for believing it. One could argue that recognising that a speaker has a justified belief is reason to take on that belief oneself: recognising that the belief is justified means recognising that there are good grounds for believing it. However, on what grounds can the hearer know that the speaker is justified in believing what she asserts? And furthermore, the hearer's presumption that the speaker's belief is justified rests on the idea that he already knows she is not lying, and actually *does* think that she knows what she asserts.

depends upon good grounds for its status as knowledge. Knowledge is a special kind of belief, with a privileged epistemic status given to it by its justification or grounds. A hearer, upon hearing a speaker make an assertion, cannot just know what she asserts; he cannot simply know that what she says is true just by understanding her. He needs to do some good epistemic work (or the epistemologist has to do it for him) in order for his resulting beliefs to gain that privileged status. That the belief is true is something that needs to be epistemologically established, and cannot just be taken for granted.

However, it is my contention that this problem, (essentially the vulnerability problem I discussed in Chapters 2 and 3), can be dissipated if we reject the internalist conception of knowledge. Knowledge is not a special kind of belief with a privileged epistemic status. Knowledge is a factive mental state. Knowledge is the mind's grasping of the world as it actually is. Knowledge is mental, it occurs within the space of reasons, and yet it is factive, it encompasses or incorporates the actual world. The mental realm is not identical to the internal or 'psychological' realm. When the speaker has knowledge that p, she does not just believe that p with good reason and p happens to be true: when she has knowledge that p, her mental state *guarantees* that what she knows is true. The possibility of false belief does not threaten or undermine the status of knowledge qua knowledge, because if she knows that p, she cannot be wrong.

I hope that all this sounds familiar, and, even better, plausible. My contention is that seeing knowledge as a mental state in this way allows for the possibility of testimonial communication to occur. The hearer can recognise the speaker's knowledge that p in the same way that he can recognise her other mental states through understanding her utterances. Successful communication entails that the hearer recognises the speaker as being in the mental state that she actually is in (the one that she is trying to get him to recognise her as being in.) In cases of testimony, that mental state is knowledge. Successful testimonial communication occurs when the hearer recognises the speaker's knowledge that p, and thus comes to know that p too. Her assertion is an attempt to get him to recognise that p in his recognition of her mental state of knowing that p, and, when all goes well, this is what he does.

The externalist conception of knowledge that I have been campaigning for is a conception of knowledge that allows the *truth* to feature as part of the mental state. Under the internalist conception, the mental components of knowledge, belief and something like

justification, are distinct from the external component: truth. For internalists, when the hearer recognises the speaker's mental state, when she knows that p, he is recognising that she believes that p, and is justified in believing that p. He cannot recognise thereby, that p is true. If all goes well, under the internalist conception, the hearer has access only to the speaker's internal state, which excludes the state of the world. The best case scenario (when a hearer understands an assertion, and takes it as sincere, and it is sincere) does not give the hearer access to the fact that p is true. This is because the truth of the speaker's assertion is not something 'within her mental grasp.' If she does know that p (that is, p is true) her assertion still cannot give the hearer access to this. All he has access to, upon understanding her assertion, is her internal mental state, which excludes the truth of p.

If the epistemologists have a good story to tell, and the hearer's resultant belief turns out to be justified, then (as justified true belief is hypothetically sufficient for knowledge) he may count as having knowledge. And perhaps, if one is fully committed to internalism, this will be the only way forward.[31] However, even if the justificatory story were good enough to establish knowledge, it still leaves an *explanatory* (as opposed to epistemological) gap. How does recognising the speaker's belief that p *get me* to believe that p? Why is it a *reason* for me to believe that p? Justification is not synonymous with good reason: for example epistemic warrant and *a priori* entitlement don't require *reasons* per se. We do not believe a speaker's testimony *because* we recognise that testimony is necessarily generally reliable, or that intelligibility is a sign of rationality, and rationality a source of truth. So why do we trust a speaker's testimony? The thin conception of communication, and an internalist conception of knowledge, I contend, cannot *explain* our behaviour of accepting testimony.

And although there will be arguments that can cover this gap too, arguments that explain how our recognition of the speaker's internal mental state can bring about knowledge of an external state of affairs, such arguments would require that our beliefs about the state of affairs are based on beliefs about the speaker's *internal* mental states. But it seems intuitively implausible to imagine that, upon hearing an assertion, I form beliefs about the speaker's beliefs, her justifications and come to the (rational?) conclusion that what she says is true. The more common sense view would be that one recognises the truth

[31] My doubts about the satisfactoriness of this route constitute Chapter Two, Section Two.

of what is asserted directly upon understanding the assertion.[32] Take a few everyday examples: The librarian asserts 'The library closes at ten'; a stranger taps me on the shoulder and asserts 'You've left your car lights on'; the petrol attendant asserts 'The price of petrol has gone up again.' In these cases, the hearer would not usually (I would argue) form any beliefs about the speaker's *internal* state whatsoever. If the librarian tells me the library closes at ten, I immediately form a belief about the library closing at ten. I do not immediately form a belief about the librarian's beliefs. I might have called the library from home, and barely registered her as a person, never mind formed beliefs about her internal mental states, beliefs that are meant to underwrite my resulting belief in the truth of what she asserts.

According to the thin conception of communication, my beliefs about the speaker's internal mental states are meant to be *epistemically prior*[33] to my belief in what she asserts, and yet I would find it strange if someone were to tell me that the only reason that I believe the library closes at ten is because I recognised that the librarian believes it closes then, and this is somehow sufficient for my knowing that it closes at ten. One could argue that the process of forming this belief happens on an unconscious level, and I suppose there is no incontrovertible reason that this should not be the case. But consider, if the belief that the library closes at ten depends, epistemically, upon the belief that the librarian believes that it closes then, and that, furthermore, she is justified in this belief, then I should be *more sure* of this than I am of the fact that the library closes at ten. And yet, in a normal case, I take myself as *knowing* that the library closes at ten, and yet, if asked, I would not straightforwardly know that the librarian's belief that the library closes at ten is justified (for example). The thought about the librarian's justifications would not even have occurred to

[32] In some interesting psychology experiments, Gilbert (1993) showed that understanding and believing an assertion seemed to be one and the same cognitive event, and that doubt occurs only retrospectively: 'Three experiments support the hypothesis that comprehension includes an initial belief in the information comprehended' (p.221).

[33] Cassam uses the notion of 'epistemically prior' in his argument against the sceptic. He claims 'Seeing that the cup is chipped is precisely a *means* of knowing, or of coming to know, something about the external world. The knowledge that it yields isn't 'epistemically prior' to the knowledge that the cup is chipped; it *is* the knowledge that the cup is chipped.' p.27. It is my belief that one can deny the epistemic priority of appearances if one embraces the externalist conception of knowledge, though Cassam does not appear to share this view, (he argues against it in 'Can the concept of knowledge be analysed?' 2009). It is interesting to note, however, that there is a parallel between Cassam's views on epistemic priority, and Williamson's (externalist) view that 'seeing that' is a 'factive mental state operator': 'if one knows that A, then there is a specific way in which one knows; one can see or remember or … that A. … We may say that knowing that A *is* seeing or remembering or … that A.' (2000 p.34).

me. On reflection, I would allow that the librarian's belief must be justified, otherwise I would not be able to know that what she asserts is true, but the reasoning seems to go from the fact that I know what she asserts to the fact that she must be justified in believing it, and not the other way around.

I do not claim that these intuitive problems are insuperable – after all, the majority of approaches to the problem of testimony are attempts at explaining (at least epistemologically) how this occurs. (The non-reductionists, for example, argue that we do not need independent grounds for trusting a speaker, but that testimony offers its own special kind of justification.) But to me these problems suggest that a more straightforward explanation would be more satisfactory. My claim is that, in recognising the speaker's intentions (that is, in understanding the assertion) the hearer is given direct access to the truth of what is asserted: a proper understanding of a communicative utterance entails recognising the speaker as being in the mental state that she actually is in, which, in the case of testimony, is the mental state of knowing. The hearer does not (necessarily) recognise anything about the speaker's *internal* state, but rather the *truth* of what she asserts. He recognises p as being true. The proper understanding of the librarian's communicative utterance that the library closes at ten is understanding (factively) *that the library closes at ten*.

Unlike the thin account, my account does not require that the hearer recognise the speaker's *internal* mental state, and on that basis come to accept that what she says is true. It is true that the hearer has to recognise the speaker's *mental state*, but this mental state is distinct in that it is factive/world-entailing. What the hearer recognises is not the internal state of the speaker, but what her mental state entails: that what she asserts is true. Access to the internal mental state of the speaker is not *epistemically prior* to having access to the fact that p. A proper understanding of the assertion gives a hearer access to the fact that p, without having to detour through the *internal* state of the speaker. Properly understanding an assertion does not require recognising *that* the speaker knows what she asserts, it requires recognising *her knowledge* of what she asserts.

My account, unlike a thin account of communication, explains *why* we accept the word of a speaker: it is because that is *how to* understand assertions. Glüer and Pagin[34] offer

[34] 1999.

an interesting discussion on the explanatory potential of constitutive rules with regards to linguistic behaviour. Their conclusion is that constitutive rules cannot explain linguistic behaviour because they are not normative. But this is because they take on Searle's '*counts as*' understanding of constitutive rules, which ignores their essentially normative nature. Constitutive rules are essentially normative, they explain behaviour because they show one *how to do* something. I do not need to explain my moving of the bishop diagonally in a game of chess with anything more than 'I know how to play chess.' The rules explain my acceptance of the librarian's assertion 'The library closes at ten' because the correct way to understand an assertion (if nothing has gone wrong) is as being true.

But what of the fact that the librarian could be lying or mistaken? If she is, does this mean that I do not *properly* understand her utterance? My answer to this is 'yes'. Something has necessarily gone wrong if the speaker breaks her rule, or if the hearer breaks his. The hearer's rule entails that he recognise the speaker as trying to get him to recognise her actual mental state, and if it is not her actual mental state that she is trying to get him to recognise, then he cannot obey his rule. This is not to say that he does not understand her *at all*, he may understand what she is trying to do, or understand what her utterance means. But, whatever understanding he has, it is not the understanding that is required in order for communication to technically succeed. Something has gone wrong when the speaker lies or is mistaken – the speaker breaks her rule, and consequently the hearer breaks his – even if it is not the hearer's fault, and even if he is unaware of the fact that he is breaking his rule.

The correct understanding of a correct assertion (testimonial utterance) entails that the hearer comes to know what is asserted. It is a specific kind of factive understanding that the hearer must have in order for that instance of communication to be a success. What this does *not* mean, however, is that the hearer *fails to understand the utterance* when the speaker is lying or mistaken, or when he takes her to be lying or mistaken. He may understand exactly what she means, and recognise perfectly what she intends to do with her utterance. He fails to understand the utterance *in the way required for communication to succeed*, he does not fail to understand the utterance *full stop*. The fact that something has gone wrong when he fails in this specific way is not so counterintuitive if one considers that this failure entails either a refusal to fulfil the speaker's intention, or a fulfilling of them that he would not have consented to if he were able to 'see' her actual mental state (it

involves understanding her to be in a mental state she is not in.). Accepting that the hearer fails to understand properly in these instances (in this technical sense) is what makes it possible for a technically *correct* understanding to give him access to the speaker's actual knowledge of what she asserts.

The rules (the hearer knowing the rules/knowing how to communicate) are what put the hearer in a position to take an assertion as true. This is the correct way to understand a correct assertion. The hearer's position is not epistemically established: he does not have to have good reasons to think that the speaker is not lying or mistaken. His position is established by the simple fact that he is a participant in the institution of communication. He knows how to understand communicative utterances, he knows that they are meant to be attempts to get him to recognise the speaker's actual mental state. Sometimes something goes wrong, sometimes there are reasons to think that something has gone wrong – there are defeaters – but when nothing goes wrong (the speaker obeys her rule, and the hearer obeys his) the hearer comes to know what the speaker is telling him.

It is not the case, as thinner accounts would have it, that the hearer needs to have *a reason* to take an assertion as true. The hearer does not need a reason to follow his rule and take an assertion as true. It would be to misunderstand my account of communication to imagine he hears and understands an assertion, and then is faced with the choice of conforming to his rule, or flouting it, and therefore requires a reason to conform to it. The rule guides *his ability to understand* the assertion. When he understands it properly he comes to recognise the truth of the assertion, this is what properly understanding the assertion entails. It is not the case that he understands it and then decides to believe it or not: his understanding it entails his believing it, it is one and the same cognitive event. And as long as nothing has gone wrong, his understanding it entails his coming to know the proposition asserted.

As long as one embraces externalism about knowledge, the threat that the possibility of the speaker lying or being mistaken poses to the hearer threatens his knowledge only in that, *if* the speaker is lying or mistaken, *then* the hearer does not get knowledge. This does not mean that the hearer is no longer in a position to get knowledge when the speaker is *not* lying or mistaken. The possibility of something going wrong does not cast a shadow over every instance of testimony, there is no epistemic gap brought about by the possibility. The hearer is able to understand an assertion as being true just by understanding it:

recognising the speaker's communicative intentions for what they are entails recognising that what she says is true.

**

The above discussion was aimed at dealing with likely objections from proponents of thinner accounts of communication. However, even those who are not involved in communication theory will have problems with my account. The rest of this chapter is devoted to offering brief responses to eight more objections. I am sure there will be many more, but I take these to pose the biggest threat to my account of how testimony-based knowledge is possible.

i. There are some consequences to my account that may sound counterintuitive. For example, my account entails that a speaker can be wrong about what she thinks she is attempting to do. The speaker *must* be in the state of knowing that p in order for her utterance to be an explicit attempt to get the hearer to recognise that p in his recognition of her knowledge that p. This means that if she does not have knowledge (though thinks that she does), her utterance is not an attempt to get the hearer to recognise that p in his recognition of her knowledge that p. She thinks she is attempting to do something (namely, get the hearer to recognise her knowledge), but (because she does not actually have knowledge) she is wrong about this. If she happens to be wrong about p, then her communicative intention is not what she takes it to be. The communicative utterance (assertion) is the speaker's attempt to get the hearer to recognise her as being in the mental state of knowing – her communicative intention is that the hearer recognise her knowledge that p. If she is not in the mental state of knowing that p, then her utterance is not an attempt to get him to recognise her as being in that state – even though she thinks that it is.

When the speaker does have knowledge, her utterance is an attempt to get the hearer to recognise her as being in a state that she actually is in. When she does not have knowledge (unbeknownst to her), her utterance is an attempt to get him to recognise her as being in a state that she is *not* actually in. What she is attempting to do in the good case is get the hearer to recognise her knowledge, but this cannot be what her utterance is an attempt to do in the bad case, because she does not have knowledge. This is important because if what she was attempting to do was the same in both cases, then what she is attempting to do in either case would not entail that she actually know what she asserts.

Consequently, the hearer's understanding of the utterance (recognition of what the speaker is attempting to do) would not entail a recognition of the fact that she actually knows what she asserts.

For speech act theorists, the illocutionary intention in cases where the speaker is lying and mistaken is the same as the illocutionary intention in cases where the speaker is being 'sincere.' But what this then forbids, is that the hearer's recognition of the speaker's intention gives him access to the speaker's actual mental state. For example, for speech act theorists, the communicative intention behind the assertion 'The capital of Peru is Lima.' is the same as the intention behind the assertion 'The capital of Peru is Havana.' The common denominator between the two intentions is that the hearer *come to believe* what is asserted. This means that, upon recognition of the speaker's intention, the hearer only has access to *the speaker's intention that he believe* that the capital of Peru is Lima (or Havana.) This leaves it open how it is that the hearer can come to know that Lima is the capital of Peru from understanding the speaker's utterance.

Under my account, however, the speaker's communicative intention is different in the case where she asserts 'The capital of Peru is Lima' to when she asserts 'The capital of Peru is Havana.' Or rather, her utterance 'The capital of Peru is Lima' counts as an attempt to get the hearer to recognise her knowledge, but her utterance 'The capital of Peru is Havana' does not count as an attempt to get the hearer to recognise her knowledge. Her intention in the first instance is to get him to recognise that the capital is Lima via his recognition of her knowledge that it is the capital. In the second instance (if she really believes that the capital of Peru is Havana) she *thinks that* her intention is to get him to recognise that the capital of Peru is Havana in his recognition of her knowledge that it is, but because she does not have knowledge, this is not her communicative intention. This may seem strange, but it is necessary if the hearer's recognition of her communicative intention is to give him access to her actual mental state. In the case where she is not mistaken, the hearer recognises that the capital of Peru is Lima via recognition of her communicative intention (which entails that she knows what she asserts.) In the case where she is mistaken, the hearer may think he recognises that the capital of Peru is Havana via recognition of her communicative intention, but he only *thinks that* he recognises it. He has taken her utterance to be an attempt to get him to recognise her as being in the state of knowing (that she actually is in),

but she is not in that state, and thus her utterance could not count as an attempt to get him to recognise her as being in that state.

On my account the speaker's utterance *is* (consists in) the attempt to get the hearer to recognise her as being in a certain mental state. When the constitutive rules are being obeyed, the utterance is an attempt to get the hearer to recognise her actual mental state, and understanding the utterance is recognising her actual mental state (in recognising her utterance for what it is, an attempt to get him to recognise her actual mental state, he recognises her actual mental state). In a sense, the mental state itself partly constitutes what the utterance is (when it is a correct utterance). Understanding her utterance properly consists in recognising the mental state (the speaker's actual mental state) that the speaker is trying to get the hearer to recognise. In the case of assertions (testimonial communication) the hearer's understanding of the utterance (his recognising of the speaker's intention – recognising what she is attempting to do) consists in the recognition of her assertion as true in recognising her (external) mental state of knowing what she asserts. This means that her communicative intention (when she is obeying the speaker's rule) entails that she knows what she asserts, and recognising this intention entails recognising that she knows what she asserts: recognising the speaker's communicative intention (understanding her utterance) entails recognising that what she says is true.

The counterintuitive claim, that a speaker can think that her utterance is an attempt to do something but be wrong about this, is explicable because the factive mental state of knowing that p *partly constitutes* what the utterance is an attempt to do. The fact that one can be wrong about what one thinks one is attempting to do is a direct fallout from the fact that one can be wrong about whether one knows something or not. Take the example of the mother who tells her son there are nine planets in the solar system. Her intention is not for her son to merely believe that there are nine planets, she (mistakenly) believes that her utterance is a means to him coming to *know* that there are. She believes that her utterance is an attempt to get him to recognise that there are nine planets via his recognition of her knowing that there are nine planets, but the fact that she does not know that there are nine planets means that she is mistaken about what her utterance is an attempt to do.

Under my account the communicative utterance is the means by which one communicates one's mental state. The consequence of this is that one is able to recognise a speaker's mental state just by understanding her utterance. The rule governing how to

make an utterance, and the rule governing how to understand one, together make it possible for communicative utterances to straight forwardly *be the means* to communicating mental states. The rule for the speaker is that her utterance *ought to be* an attempts to get the hearer to recognise her actual mental state. But a consequence of this rule, and the fact that it is possible to think that one is in the mental state of knowing when one is not, is that a speaker can think that she is following her rule (that her utterance is an attempt to get the hearer to recognise her actual mental state) and yet not be. She can think that her utterance is an attempt to do something that it is not.

ii. Another counterintuitive consequence of the fact that knowledge is a state that one can think one is in, and yet not be in, is that the speaker can attempt to obey her constitutive rule, and fail to do so. This may seem counterintuitive because the rule demands not only that the speaker do her epistemic best in forming a belief, but that that belief be *true* as well. If the speaker has done everything to ensure that her belief is true, how can she be doing something wrong, if, by bad luck, her belief turns out to be false? Is it not true that *ought implies can*? How can the speaker's rule demand more of her than that she do her epistemic best?

However, this objection presumes that there is no mental difference between a speaker who has done her epistemic best and knows, and a speaker who has done her epistemic best and fails to know. It is assumed that whether or not one knows is something that is determined purely by the external world. And this, of course, is denied by externalism. The externalist view holds that knowledge is a different mental state to justified belief. Justified belief that is false is a *failed* belief. The speaker has 'done something wrong' if her justified belief is not knowledge. And when she does have knowledge, she has 'done something right.' If you allow luck to 'enter the space of reasons' as McDowell advocates, then it is possible to see how knowledge is the epistemic achievement that justified belief is not. When one knows something, the world is as one takes it to be, and not because the world happens to be the way that one takes it to be, but because the *mental* state of knowing it entails that it is the way one takes it to be.

The upshot of this is that the speaker's rule does not prescribe an ought where there is no 'can.' If the speaker can know what she asserts, then she can obey the rule. And being in the mental state of knowing is a state that she achieves as a rational agent. It is true that

she may fail to know when she thinks that she does know, but this does not mean that knowledge is not a mental state, it just means that she can sometimes be wrong about what mental state she is in. The rule does not demand more of the speaker than she is capable of. She can obey the rule by asserting only what she knows. The fact that she can sometimes think she is obeying the rule, and yet not be, reveals that knowledge is not a luminous mental state, and nothing more.[35]

The above objection implies that the rule ought to be that the speaker be justified in believing what she asserts, seeing as that is (supposedly) the 'best epistemic position' that one can achieve. But the fact that knowledge is the rule and not justified belief is evident in the fact that if one discovers that one has asserted a falsehood by accident, one would view that action as incorrect. If the librarian suddenly remembers that it is Saturday, and that the library closes at nine and not ten, she would regard her earlier assertion to me, 'The library closes at ten', as incorrect. And not because she regards her belief as less justified, but because she realises that she did not know what she purported to know when she responded to my question. She told me something that was false. If the rule was that she ought to be justified in her belief, she could judge her assertion as still correct (after all, her belief that is was Friday was based on reading the calendar, and was thus justified.)[36]

iii. Another possible point of resistance to my account may be to the claim that the hearer is doing something wrong by rejecting a false assertion. If the son has recently learnt in school that there are actually only eight planets in the solar system, he may reject his mother's assertion that there are nine. From an epistemological point of view, the son is doing the right thing. He has reasons to think that his mother does not know what she asserts, and therefore he ought not to take her word for it. To call his rejection wrong, incorrect, or breaking the rules, seems counterintuitive. He should not have done otherwise.[37]

---

[35] Timothy Williamson (2000) argues against luminosity of mental states in Chapter 4 of *Knowledge and its Limits.*

[36] Williamson offers good arguments against accounts of assertion that invoke a rule about belief or justified belief rather than knowledge (2000, Chapter 11). Douven 2006, Kvanvig 2009, and Lackey 2007 argue for a justified belief rule (or something similar).

[37] It might be worth noting that if the speaker *is* mistaken, the hearer breaks his rule when he accepts her utterance too – not only if he rejects it. (If the speaker breaks her rule, the hearer *cannot* obey his.) But this is not because of epistemic reasons, but because his rule states that he has to take her utterance as an attempt to get him to recognise her actual mental state, and if she is not doing this, then he cannot properly recognise her as doing this.

With this I agree. A hearer who has reason to think that the speaker is lying or mistaken ought not to take her at her word. He ought not to believe what she says. If he does believe her despite reasons not to, he is being particularly gullible or credulous. However, what this illustrates is not the weakness of the constitutive rule, or that it does not exist. What it illustrates is that the *kind* of normativity imposed by constitutive rules is not a moral or epistemic 'ought.' These rules form the foundations of communication, they are 'how to' rules rather than 'have to' rules. They distinguish between correct and incorrect communicative behaviour, but this is the full extent of their power. They explain how it is that communication works, but they do not have (an intrinsic) moral or epistemic (i.e. justificatory) dimension. The constitutive rules of communication explain how it is that knowledge can be transmitted through communication, not how the beliefs we get from testimony are 'secure' or justified enough to count as knowledge. (The externalist account of knowledge denies that being sufficiently justified is a necessary condition on knowledge).

If one is asked 'How do you know that the price of petrol has gone up?', the answer 'I was told so' is a perfectly acceptable answer. The answer does not explain why believing that the price has gone up is a *good* belief, or how it has the 'privileged epistemic status' of knowledge, but it does explain *how it is that I came to know it*. It might become a less acceptable answer when there are reasons to think that something might have gone wrong. Perhaps the questioner has been told that the petrol price has gone down. In this situation, the two could discuss where they got their information, and attempt to work out which claim is most likely to be true. However, it is only when there are other, defeating epistemic reasons in play that 'I was told so' is no longer a sufficient response to the question 'How do you know?' And if one's belief is challenged by defeaters, one will need more than just 'I was told so' to maintain that one knows what one was told. 'I was told that the capital of Peru is Havana' is not enough to defend one's epistemic position in the face of evidence to the contrary.

An analogy can be drawn between the constitutive rules of communication, and the enabling conditions Cassam proposed to explain how perceptual knowledge is possible. Cassam proposes two *a priori* enabling conditions for the possibility of perceptual knowledge: the CTR (Categorical Thinking Requirement) and the SPR (Spatial Perception Requirement). These conditions he considers necessary to explain how we can get knowledge through 'seeing that' something is the case, and yet these conditions do not play

a justificatory role, they do not show how the beliefs we get through 'seeing that' something is the case have the distinguished epistemic status of knowledge as opposed to mere belief. They are not obstacle overcoming conditions, they are necessary conditions that can *explain how* knowledge can be acquired in this way. The question 'How do you know Bob was at the party?' can be acceptably answered with 'I saw him.' The enabling conditions explain how this is possible. Once there is reason to doubt this, to think that something has gone wrong, perhaps a justification will be necessary to defend the fact that one knows that Bob was at the party.

The institution of communication, the constitutive rules that make communication possible, play this role. They explain how knowledge can be gained through understanding an assertion without being justifications for belief. Hearing it asserted that p is *a way of coming to know* that p because the institution of communication is a rule governed activity, and knowing the rules puts one in a position to know that p just from understanding the assertion that p. The rules explain the possibility of coming to know in this way without (necessarily) having moral or justificatory dimensions. The rule that the hearer ought to understand an utterance as an attempt to get him to recognise the speaker's actual mental state is not a rule that *justifies* the hearer's acceptance of the speaker's word: it explains it.

The rule that one ought to take an utterance as an attempt to get one to recognise the speaker's actual mental state is a rule that can be defeated by other considerations. Of course, one must use one's reason to distinguish between what count as defeaters, and what do not. The hearer, in order to count as a knower at all, must be responsive to reasons that may threaten or undercut his knowledge, and the same is true of getting knowledge through perception. But if there are no reasons to think anything has gone wrong, and, furthermore, nothing *has* gone wrong, then one can get knowledge of what is asserted. But, the hearer does *not* need reasons (moral or epistemic) to think that everything has gone right, this is just how he ought to understand assertions, all things being equal.

The fact that it is often considered a 'bad thing' if the speaker breaks her rule (by lying or just being wrong) is *not* a result of the constitutive rule of communication, it is because there are other moral rules guiding the speaker's behaviour. Lying is considered bad for reasons other than the fact that the possibility of communication defines it as incorrect. This can be seen by the fact that it is not *always* bad to lie. Lying is often considered morally wrong, perhaps because it involves deception. But sometimes there are other moral

considerations, such as in the famous 'hiding Jewish refugees from the Nazis' thought experiment. In this situation (the Nazis are at the door, demanding information) it is morally preferable to lie to the Nazis than to reveal the refugees' hiding place. But even though it is morally correct to lie in this situation, it is an instance of flawed or technically unsuccessful communication: it is incorrect because the speaker tries to get the hearer to recognise her as being in the state of knowing that there are no Jewish refugees in her home, but she does not know this. (And her lie can only succeed in the way she intends if the Nazis take her assertion to be attempt to get them to recognise her *knowledge* – and thus that it is true – that there are no Jewish refugees in her home – in other words, the fact that it is only correct to assert what one knows is what makes it possible for her lie to be taken as sincere.)

In consequence, there is no conflict between the claim that it is a constitutive rule of communication that the hearer ought to take the speaker as trying to get him to recognise her as having or being in the mental state that she actually is in, and that one is often doing the right thing by *rejecting* an assertion as such. There is no conflict because the obeying of the constitutive rule is not the right thing, *all things considered*. Often there will be epistemic or moral considerations or defeaters that make rejecting an assertion the right (epistemic or moral) thing to do. The breaking of the constitutive rule is not a *bad* thing to do, it is just an incorrect thing to do within the institution of communication.

iv.) A related objection follows naturally on from this. If the constitutive rule for the hearer does not justify his resultant belief, then surely the rule promotes gullibility. The hearer accepts an assertion as an attempt to get him to recognise the speaker's knowledge, but his acceptance is not *justified*. Taking a speaker's word for it is not something that is *rationally* defensible, it is just what one ought to do given that one is a participant in the institution of communication. Gullible people are people that believe things that they are not justified in believing, and if the constitutive rule guiding the hearer's interpretation of utterances is not a justification, then accepting what the speaker says is a gullible thing to do.

Firstly, this objection arises out of a commitment to the internalist conception of knowledge, a view that makes it a necessary condition on knowledge that the belief is justified. But, as I have argued, we ought not to think of knowledge in these terms. Consider what we actually mean when we use the term gullible. Gullible people are not judged

gullible because they believe things that they are *not justified* in believing. Gullible people tend, more than non-gullible people, to believe things that there *are reasons to doubt*. Gullible people are bad at spotting defeaters. They come to have false beliefs more easily than non-gullible people. To claim that it is gullible to believe something without *positive justification* or proper epistemic warrant/entitlement is not accurate use of the term. Consider the fable of the boy who cried wolf. The boy comes down to the village and cries 'Wolf!' a few times but when the villagers run to assist, there is no wolf. A person who continues to believe the boy after he has cried wolf on many occasions would be considered gullible. He still believes the boy *despite there being a reason to doubt* his testimony. A visitor to the village who does not know the boy's unreliable history is not gullible for believing the boy the first time she hears him cry 'Wolf!' She is in a different epistemic position to those who are aware of the boy's untrustworthiness. She has no positive reason (or entitlement) to believe him, and yet accepting his word is not *gullible*.[38]

When there are no reasons to doubt the truth of an assertion, when there is no reason to think that something has gone wrong, a gullible person (a person who is bad at spotting defeaters) and a non-gullible person are *equally* in a position to gain knowledge from understanding the assertion. If nothing has gone wrong, and there are no defeaters, the correct way to understand a correct assertion entails coming to know what is asserted. The fact that the constitutive rule for the hearer does not provide justification or epistemic warrant does not undermine the hearer's ability to get knowledge thereby. Rationality or responsiveness to reason is a necessary background to the possibility of knowledge (as I made clear in Chapter 3), but unless one is non-rational, one will still be able to get knowledge from testimony. One does not require justification or a story about reliability or epistemic entitlement in order to explain how knowledge can be gained from testimony.

v. Another objection to my account could arise out of a rejection of my use of the word 'understanding'. I claim that properly understanding an assertion requires or entails coming to know what is asserted. This claim goes against the intuitive grain in that it implies that one has not understood an assertion properly if one fails to accept it as true. It seems clear that I can understand an assertion such as 'I was abducted by aliens,' perfectly, without

[38] McDowell uses the cry wolf story in 'Knowledge by Hearsay' (1994) to illustrate his account of doxastic responsibility, p.211-212.

accepting what has been said. I do not *fail to understand* in any intuitive sense. Correlatively, accepting a false utterance as true does not seem to be a case of failing to understand either. If I believe the speaker who tells me he was abducted by aliens, it seems wrong to say that I did not *properly understand* him. Claiming, as I do, that properly understanding a testimonial utterance entails coming to know what is asserted, is to misapply the term 'understanding.'

However, the term 'understanding' is broad and complex, and the sense in which I apply it to communication is only *one specific sense* of understanding. When I claim that proper understanding entails coming to know what is asserted, I mean proper understanding as delineated by my account of communication. Proper understanding in the sense required if communication is a success requires recognising the speaker's actual mental state, and that state is the one that she is attempting to get the hearer to recognise through her utterance. Understanding still occurs when one recognises the speaker's actual mental state, and yet it is not the mental state she is trying to communicate: this is not proper understanding under my account. Understanding still occurs when one understands what mental state the speaker is trying to communicate, and yet does not see it as being the speaker's actual mental state: this is not proper understanding on my account. Understanding an utterance, broadly construed, can have many implications. My account of communication is to delineate a *special kind* of understanding that occurs only when communication is successful: this occurs only when both speaker and hearer are conforming to their rules. The fact that this specific notion of understanding is not always achieved when we would intuitively claim understanding, *broadly construed*, has been achieved, does not threaten my account.

One could argue that the objection obliges me to use some term other than 'understanding'. And yet, I do mean 'understanding'. I mean it in the sense that an innocent person who claims 'I am innocent, don't you understand?' means it. She is not saying, 'Do you understand what my utterance means?' She is not saying, 'Do you understand that I want you to think that I am innocent?' or 'Do you understand that I believe that I am innocent?' She is saying 'Do you understand that I AM innocent.' If the hearer does not take her to be innocent, he has not *understood* in the sense that she is imploring him to. The sense of understanding I employ is the sense in which a speaker means it when she says: 'Class starts at two o'clock. Do you understand?' The hearer has to

take on the knowledge proffered, has to know what the speaker knows, in order to understand in this sense.

Broader conceptions of the word 'understanding' do not capture all that can be meant by the word. The kind of understanding that my account of communication requires is the understanding that comes along with acceptance of what is asserted. The idea that understanding an assertion and accepting it as true (believing it) can be one cognitive event is supported by psychology experiments.[39] In fact, the experiments seemed to show that not only is it possible for understanding an assertion and accepting it to be one cognitive act, but that it *always* is. When we disbelieve an assertion, this occurs only after it has been initially understood and *believed.* I do not need to make such a strong claim however (though the results stand in my favour). For my purposes, all that needs to be plausible is that there is an important notion of understanding an assertion that entails recognising what is asserted *as true*.

Understanding utterances is usually taken to mean understanding *what the utterance means*, or *what the speaker is intending with the utterance*. Outside the realm of communication theory, however, understanding is often (usually) taken to be factive: one's theory that explains the movement of the stars does not count as *understanding* the movement of the stars, unless the theory is correct. Understanding what has happened requires that one *knows* what has happened. It would be wrong of me to claim that I understand how motor engines work if my understanding does not entail *knowing* how engines work. (If I said that I thought engines ran on fairy dust I could not claim to *understand* how engines work.) I want to import this factive notion of understanding into communication. A correct testimonial communicative utterance is an attempt to get the hearer to recognise that p in his recognition of the speaker's knowledge that p. Understanding it properly entails recognising it for *what it is*. It is not an attempt to get the hearer to believe that p, it is not an expression of a belief that p, it is not an utterance that may or may not reflect the speaker's actual mental state. It is an utterance that is partly constituted by the speaker's actual mental state of knowing that p, and recognising it for what it is entails recognising that p is true.

[39] Gilbert et al. (1993).

vi. Another point of resistance to my account may be the existence of some problematic counter examples. There are some instances of communication that are prima facie assertions and prima facie successful (under my account), and yet do not entail the hearer coming to know what is asserted, for example my brother asserts 'Broccoli is disgusting' and I understand him perfectly (recognise his actual mental state) without coming to know that broccoli is disgusting. Or a religious protester asserts 'Same sex marriage is wrong' and the passer-by who understands her will have been successfully communicated with. The speaker's utterance is an attempt to get the hearer to recognise her actual mental state, and the hearer thereby comes to recognise her actual mental state, without thereby coming to know what the speaker asserts.

However, I would contest that these situations are not instances of testimony. I allowed (n.7 of this chapter) that assertions are not always testimonial utterances, and these situations would be examples of such assertions. When my brother asserts 'Broccoli is disgusting' he is not *telling* me that broccoli is disgusting, he is telling me that he hates broccoli. He is expressing his opinion, and he does so successfully if I come to recognise that this is his opinion. The same is true to the 'Same sex marriage is wrong' assertion. The protester is not *telling me* that it is wrong (which would imply that I had no previously held view on the matter, and her testimony would be the way to enlighten me), she is trying to *convince* me of something. She hopes that her opinion will be a reason for me to change mine. It would be an inappropriate response for me to say 'Is it? Oh, I did not realise!' The protester's intention is not to straightforwardly tell me of the moral wrongness of same sex marriage.

The key difference between these cases and testimony is that the mental state that the speaker is trying to communicate is not knowledge. My brother does not know that broccoli is disgusting: if he did, it would be an objective state of the world that broccoli was disgusting. He knows that he does not like it, and it is this that he successfully communicates to me. The protester may think that she knows that same sex marriage is wrong, but she does not. And not because same sex marriage is *not* wrong (though I believe that to be the case) but because moral judgements do not have a truth value in the same way that states of affairs do. This, of course, will lead me into very dangerous territory, territory that I do not wish to explore here. However, it is worth pointing out that if a hearer did come to believe that same sex marriage was wrong after hearing the protester, it

would be strange if his response to the question 'Why do you think same sex marriage is wrong?' was 'Because that woman told me so.' His belief in the immorality of same sex marriage is not *testimony-based knowledge*. (Not in the sense that I am interested in explaining, anyway.)

Assertions are not always instances of testimony, and we can tell which ones aren't by looking at the speaker's communicative intention. If she is attempting to persuade me, convince me or encourage me to believe something, it is not testimony. Testimony occurs when the speaker's intention is to *tell* me something. Her expectation must be that simply upon understanding her assertion I will come to know it is true. The stranger who asserts 'You left your car lights on' is not attempting to persuade me of a certain state of affairs, she is attempting to *let me know*. It is her intention that I come to know that I left my car lights on just by hearing and understanding what she has said. The same cannot be said of my brother's intention behind 'Broccoli is disgusting' or the protester's 'Same sex marriage is wrong.'

vii. One of the necessary conditions for successful testimonial communication in my account is that the speaker knows what she asserts, as this is the means by which she communicates her knowledge to the hearer. Jennifer Lackey, in her book, *Learning from Words*[40] argues that testimony can still succeed, even if the speaker does not know what she asserts. She gives an example of a devout Creationist, who also happens to be a biology teacher.[41] The teacher teaches evolution to her students even though she does not even believe (never mind know) what she tells them. Lackey's point here is that a focus on knowledge *transmission* in testimony is misplaced: knowledge can be *created* through testimony. The consequences of this, for Lackey, are far reaching, but I believe her example to be flawed. The hearer's knowledge *must* come from knowledge. The example of the Creationist teacher, however, illustrates an interesting dimension of testimony.

The teacher, when she teaches evolution, is not straightforwardly telling her students that modern humans evolved from apes. She is an *instrument* of an educational institution, a mouthpiece, as it were. Her authority comes not from her knowledge as an individual, but by her position as a teacher. She is the vehicle *through which* knowledge is transmitted to

[40] 2008.
[41] 2008 p.48.

students. The students can come to know that humans evolved from apes because she holds the special position as teacher: her job is to make that knowledge available to them. In Lackey's example, the teacher 'regards her duty as a teacher to involve presenting material that is best supported by the available evidence.'[42] I believe that it is her duty as a teacher to teach *what is known*, and it is only in this context that the students can come to know what she asserts. If asked about the origins of humanity, outside of the school context, she would probably reply that we were made in God's image and do not descend from apes. But her job, her role as teacher, is to disseminate what is known to her students, and this is what she does when she teaches evolution (even if she believes it to be false.)

In the same way, we get knowledge from institutions such as newspapers and TV. The 'speaker' is not the presenter, we are not believing Mark, or Elizabeth the newsreader when she tells us that there has been another aeroplane disaster. She is just the means by which the news is transmitted to world. We would not call her a liar, or hold her accountable if what she reported happened to be false – we would blame the TV channel or the newspaper.[43] She does not even have to believe what she reports, but this makes no difference to the fact that we can come to know what 'she' tells us. Usually the teacher or the presenter will know what she asserts, but if she does not, it need not be fatal to the possibility of their assertions communicating knowledge. However, if what they assert is not known, even if believed to be true by all involved, this *is* fatal to the possibility of their assertions communicating knowledge.

So who is the speaker in these cases? Whose mental state of knowing are we recognising when we recognise that what is asserted is true? There does not need to be a specific or individual speaker. This might, at first, sound problematic for me, seeing as my account holds that the hearer recognises that what is asserted is true in recognising the speaker's mental state of knowing it. But as long as we are not committed to the notorious internalist conception of knowledge, this is not as problematic as it sounds. The mental state of knowledge is not internal, it is an 'external' or factive mental state. The hearer is not recognising the *internal* state of a speaker, she is recognising an 'external' mental state –

[42] 2008 p.48.

[43] This may not be true for 'on location' journalists. Journalists reporting from the scene of the event in question would be understood to be individuals telling the news team, and the rest of the world, what is happening. In this way they are different from news presenters reading a report, or (depending on the context) teachers teaching a syllabus.

recognising, upon understanding 'p', *that p*. As long as what is asserted by the presenter *is known* (even if not by her), this is enough for knowledge to be transmitted. The teacher and the news presenter report what *is known* and this is sufficient for the hearer to get knowledge. The teacher and the news presenter are instruments of communication, their job is to report what is known, and it is in this capacity that their assertions are understood.[44]

I am sure that there are exceptions to this, and that sometimes, depending on how the example is described, the teacher not knowing what she teaches *would* be fatal to the possibility of her students coming to know it. However, in the case of Lackey's Creationist teacher, and in other cases where we would allow that the hearer gets knowledge from a 'speaker' who does not, I contend that there must *be* knowledge that the hearer is recognising in order for the hearers to get it. [45] The hearer's knowledge of what is asserted depends upon the knowledge of the speaker, it does not come just from, as Lackey contends, 'what the speaker says.'[46] Communication of knowledge entails that the speaker (the one(s) whose knowledge is being communicated) has knowledge, otherwise the hearer could not get it from understanding the assertion. Lackey's case reveals that the hearer might not even know whose knowledge he is recognising, but this does not matter if what he is recognising is its factiveness. Understanding a teacher's assertions, assertions she makes qua teacher, entails recognising what *is known.*

Successful communication requires that the hearer recognise the speaker as being in the mental state that she is trying to get him to recognise her as being in, and that this is her actual mental state. For testimony, that mental state is knowledge, which is a distinctive type of mental state: it is factive/world entailing. It is this distinctive external feature that is communicated through testimony, and this distinctive feature that allows for the speaker's intention to be to communicate *that p*, not her internal mental state. And it is also this distinctive feature of knowledge that allows for the hearer to be successfully communicated with, even if he does not know who the asserter is. As long as what is asserted is known,

[44] Similar stories could be told for how a messenger, who does not understand the message, is still able to communicate it to a hearer. What is communicated must *be known*, and this is what is recognised by the hearer, not who the knower is or even that there is a knower.
[45] For a more thorough discussion on Lackey's challenge, see Carter 2014.
[46] 2008, Chapter 3.

and the hearer is in a position to recognise that it is known, he can get knowledge from it just by understanding the assertion.

viii. One final objection that I would like to consider is an objection to the overall satisfactoriness of the enabling condition of communication in explaining the possibility of testimony-based knowledge. As I have mentioned (Chapter 4, n.1), the enabling condition I set out to identify, namely, the institution of communication, is not the *only* necessary condition for the possibility of testimony-based knowledge. Other conditions must also be in place in order for knowledge to be transmitted with an assertion. One that may be of particular worry is the condition of *the ability to use a language*. In order for the hearer to understand a speaker's assertion and thus get knowledge from him, it is not enough for him to know the constitutive rules of communication. He must know the meaning of the words she has used, he must be able to recognise the grammar, the tone she uses, and other contextual clues in order to determine *that she is making an assertion*, and not making a joke, being sarcastic, or singing a song.

Without this kind of ability, the hearer's knowing the rule that he ought to take the utterance as an attempt to get him to recognise the speaker's actual mental state will not explain how he gets knowledge just from an instance of testimony. This objection has force. In order to *fully* explain how knowledge can be gained from testimony, much must be said about the ability to understand language and conversational cues, etc. However, it is my hope that I can get away with saying that *if* a hearer is able to understand an assertion (and this will depend on a few enabling background conditions being in place) then he can get knowledge from understanding it, because his understanding it properly entails his coming to know what is asserted. It may be an important job to set out what all these conditions are, but it is a job for another day.

Conclusion

The possibility of testimony-based knowledge is explained by the institution of communication: the grounding rules of communication, the constitutive rules, make it possible for a speaker to communicate to a hearer her mental state. In the event of testimony, that mental state is knowledge. The speaker, by making an assertion correctly

(as an attempt to get the hearer to recognise that p in recognising her (actual) mental state of knowing that p) is able to communicate *that p* to a hearer. The hearer, in understanding the assertion that p is in a position to recognise that p *in* recognising the speaker's knowing that p: properly understanding the assertion that p *is* coming to know that p. The speaker's mental state of knowing that p entails that p, and it is this factiveness (the fact that p) that the hearer is recognising upon hearing her assertion. A successful instance of testimonial communication entails that the hearer comes to know what is asserted.

The possibility of the hearer coming to know what the speaker asserts is not underwritten by epistemic reasons or justification. His being a participant within the institution of communication (which almost all humans are) puts him in a position to get knowledge in this way. Things can go wrong, and the hearer, as a rational agent, will be aware of this possibility. However, things can go wrong without his knowing, and things can go wrong without him or the speaker knowing. But this does not threaten the possibility of getting knowledge in this way, because the rules stand *necessarily*. If he could not get knowledge in this way, it would be doubtful that he knew how to communicate. Knowing how to communicate means knowing the rules, knowing the rules means being in a position to take an assertion 'p' to mean *that p*. Knowing how to communicate means being in a position to know that Lima is the capital of Peru just from hearing and understanding the assertion 'Lima is the capital of Peru'.

At the end of the third chapter of this thesis, I claimed that in order to give a satisfactory explanation to the question of how is testimony-based knowledge possible (the third level of Cassam's structured response), I must be able to show that there is no epistemic gap to be crossed between hearing an assertion that p and coming to know that p. This is what I hope to have established (or at least made plausible) in the last two chapters. The problem of testimony is not a problem of how we are justified in accepting the word of another given the necessary indirectness of testimony as a source: it is not a problem of epistemic vulnerability. We are not faced with an epistemic gap between hearing an assertion and coming to know what is asserted, an epistemic gap that can only be crossed with justification. There is no gap. Hearing and understanding an assertion, all going well, *is* coming to know what is asserted. The aim of these last two chapters has been to provide a substantive explanation for this, to establish how it is that testimony-based knowledge can be had.

**

I have often found that there are moments in philosophy where one's reasoning forces one to give up on one's closely held intuitions. Sometimes we ought to follow our reasoning, no matter where it takes us, but sometimes it might be better to question the line of reasoning than to give up those intuitions. It is my closely held intuition that getting knowledge from testimony is a straight forward and mundane occurrence, and I believe that a line of reasoning that would force me to give up on that is a line of reasoning that we ought to question. If my conception of knowledge were such that knowledge from testimony became difficult to defend, vulnerable, or undermined by the possibility of error, then it would be a conception of knowledge that I ought to question. If my conception of communication were such that it did not explain how utterances communicated states of affairs, it would be a conception of communication that I would have to question.

In this thesis I have espoused what might be considered an unpopular conception of knowledge, and proposed what might turn out to be an unpopular conception of communication. But in so doing, I hope that I have been able to defend the idea that knowledge through testimony is a mundane and straight forward occurrence. Knowledge is not a kind of belief that requires a special or especially strong justification in order to establish that it is knowledge. We do not need to establish that testimony offers a special or especially strong kind of justification for belief in order to establish that we can get knowledge in this way. Knowledge is just what we have when we have a cognitive grasp on the way the world is – the world, as it is, is directly available to us. And when we get knowledge through testimony, we have a cognitive grasp on the way the world is, and the possibility of the speaker lying or being mistaken does not undermine the possibility of this occurring.

And testimony can be a way of transmitting knowledge because communication is not about performing speech acts that may or may not reflect the speaker's actual mental state. Communication is the means by which we (as speakers) get hearers to recognise our actual mental states, and the means by which we (as hearers) can recognise the mental states of speakers. And when we use assertions (testimonial communicative utterances) we get the hearer to recognise our knowledge of what is asserted. When a hearer properly understands an assertion, he comes to know what is asserted. Communication is the means

by which knowledge is shared. We do not need to rely on extra-communicative conditions in order to explain how it is that hearing an assertion that p can get a hearer to know that p. Successful communication entails that the hearer come to know that p upon understanding an assertion that p.

The aim of this thesis was to give a response to the question of how testimony-based knowledge is possible. The question arises as a result of a perceived threat to knowledge, namely the threat of epistemic vulnerability posed by the second-hand nature of testimony. The fact that knowledge through testimony is based on the knowledge of a speaker, and not gained via one's own perception or reason, is seen to pose a threat to the possibility of getting knowledge in this way: a speaker could always be lying or mistaken. However, I argued that we could ignore this threat if we took on the externalist conception of knowledge, and that we should focus, instead, on explaining *how* an assertion can function as a means to transmitting knowledge. The answer to the question of how testimony-based knowledge is possible is that we can just get knowledge from hearing an assertion because we know how to communicate.

***

Bibliography

Adler, J. (2002) *Beliefs Own Ethics*. MIT Press, Massachusetts.

Adler, J. (2014) 'Epistemological Problems of Testimony' in *The Stanford Encyclopaedia of Philosophy* (Spring Edition) (ed. E.N. Zalta). URL = http://plato.stanford.edu/archives/ spr2014/entries/ testimony-episprob/.

Austin, J. L. (1962) *How to Do Things with Words*. Harvard University Press, Cambridge, Massachusetts.

Bach, K. (1981) 'Escaping the Speech Act Dilemma', *Analysis* 41(3): 146-149.

Bach, K. (1994) 'Conversational impliciture', *Mind & Language* 9(2): 124-62.

Bach, K. 'Speech Acts' in *Routledge Encyclopaedia of Philosophy*, URL = http://online.sfsu.edu/ kbach/spchacts.html.

Bach, K. and Harnish, R. M. (1979) *Linguistic Communication and Speech Acts.* MIT Press, Massachusetts.

Ball, B. (2014) 'Deriving the Norm of Assertion', *Journal of Philosophical Research* 39: 75-85.

Benton, M. (2011) 'Two More for the Knowledge Account of Assertion', *Analysis* 71: 684-687.

Benton, M. (2012) *Knowledge Norms: Assertion, Action, and Belief*, unpublished dissertation. URL= http://mss3.libraries.rutgers.edu/dlr/showfed.php?pid=rutgers-lib:37370/.

Brandom, R. (1994) *Making it Explicit.* Harvard University Press, Cambridge Massachusetts.

Breuckner, A. (2005) 'Fallibilism, Underdetermination, and Skepticism', *Philosophy and Phenomenological Research*, 71(2): 384-391.

Brown, J. (2005) 'Williamson on Luminosity and Contextualism', *The Philosophical Quarterly,* 55(219): 319-327.

Brown, J. (2008) 'The Knowledge Norm for Assertion', *Philosophical Issues* 18, Interdisciplinary Core Philosophy: 89-103.

Brown, J. & Capellen, H. (eds.) (2011) *Assertion: New Philosophical Essays.* Oxford University Press, Oxford.

Burge, T. (1993) 'Content Preservation', *The Philosophical Review* 102(4): 457-488.

Burge, T. (1996) 'Interlocution, Perception and Memory' in *Philosophical Studies* 86(1): 21-47.

Carter, J.A. & Nickel. P.J. (2014) 'On Testimony and Transmission', *Episteme* 11(2): 145-155.

Cassam, Q. (2007) *The Possibility of Knowledge.* Oxford University Press, Oxford.

Cassam, Q. (2009) 'Can the Concept of knowledge be Analysed?' in *Williamson on Knowledge* (ed. P. Greenough and D. Pritchard). Oxford University Press, Oxford.

Cherry, C. (1973) 'Regulative Rules and Constitutive Rules', *The Philosophical Quarterly*, 23(93): 301-315.

Clark, H.H. (1996) *Using Language.* Cambridge University Press, Cambridge.

Coady, C.A.J. (1992) *Testimony: A Philosophical Study*. Oxford University Press, Oxford.

Cohen, S. (1988) 'How to be a Fallibilist', *Philosophical Perspectives* 2, Epistemology: 91-123.

Cottingham, J. (ed., trans.) (1996) *Descartes: Meditations on First Philosophy* (Second Edition). Cambridge University Press, Cambridge.

Davis, S. (1980) 'Perlocutions', in *Speech Act Theory and Pragmatics* (ed. J.Searle, F. Kiefer & M. Bierwisch). D. Reidel, Dordrecht.

DeRose, K. (2002) 'Assertion, Knowledge and Context', *The Philosophical Review* 111(2): 167-203.

Descartes, R. (see Cottingham).

Dougherty, T. & Rysiew, P. (2009) 'Fallibilism, Epistemic Possibility, and Concessive Knowledge Attributions', *Philosophy and Phenomenological Research* 78(1): 123-132.

Douven, I. (2006) 'Assertion, Knowledge and Rational Credibility', *The Philosophical Review* 115(4): 449-485.

Faulkner, P. (2006) 'Understanding Knowledge Transmission', *Ratio* 19(2): 156–175.

Faulkner, P. (2007) 'On Telling and Trusting', *Mind* 116(464): 875-902.

Fisher, J. (2007) 'Why Nothing Mental is just in the Head', *Noûs* 41(2): 318-334.

Fricker, E. (1987) 'The Epistemology of Testimony', *Proceedings of the Aristotelian Society, Supplementary Volumes* 61: 57-83.

Fricker, E. (1994) 'Against Gullibility', in *Knowing from Words* (ed. B.K. Malital & A. Chakrabarti). Kluwer Academic Publishers, Dordrecht.

Fricker, E. (2006) 'Second-hand knowledge', *Philosophy and Phenomenological Research* 73(3): 592-618.

Fricker, E. (2009) 'Is Knowing a state of Mind? The Case Against', in *Williamson on Knowledge* (ed. P. Greenough & D. Pritchard). Oxford University Press, Oxford.

Fricker, M. (2007) *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford University Press, Oxford.

Gettier, E. (1963) 'Is Justified True Belief Knowledge?', in *Analysis* 23: 121–123.

Gilbert, D.T., Tafarodi, R.W., & Malone, P.S. (1993) 'You can't not believe everything you read', *Journal of Personality and Social Psychology* 65(2): 221-233.

Glüer, K. & Pagin, P. (1999) 'Rules of Meaning and Practical Reasoning', *Synthese* 117(2): 207-227.

Glüer, K. & Wikforss, Å. (2009) 'The Normativity of Meaning and Content', in *The Stanford Encyclopedia of Philosophy* (Summer Edition) (ed. E. N. Zalta). URL = http://plato.stanford.edu/archives/ sum2009/entries/meaning-normativity/.

Goldberg, S. (2006) 'Reductionism and the Distinctiveness of Testimonial Knowledge', in *The Epistemology of Testimony* (ed. J. Lackey & E Sosa). Oxford University Press, Oxford.

Goldberg, S. (2007) *Anti-individualism: Mind and Language, Knowledge and Justification*. Cambridge University Press, Cambridge.

Goldberg, S. (2009) 'The Knowledge Account of Assertion and the Nature of Testimonial Knowledge', in *Williamson on Knowledge* (ed. P.Greenough & D. Pritchard). Oxford University Press, Oxford.

Goldberg, S. (2011) 'Assertion and the Ethics of Belief', David Klein Philosophy Symposium. URL = https://www.academia.edu/3544545/Assertion_and_the_Ethics_of_Belief.

Goldberg, S. (2011) 'Putting the Norm of Assertion to Work: The case of testimony', in *Assertion: New Philosophical Essays* (ed. J. Brown & H. Cappelen). Oxford University Press, Oxford.

Goldman, A. (1967) 'A Causal Theory of Knowing', *The Journal of Philosophy* 64(12): 357-372.

Goldman, A. (1999) 'Internalism Exposed', *The Journal of Philosophy* 96(6): 271-293.

Goldman, A. (2011) 'Reliabilism', in *The Stanford Encyclopedia of Philosophy* (Spring Edition) (ed. E.N. Zalta). URL = http://plato.stanford.edu/archives/spr2011/entries/reliabilism/.

Graham, P.J. (2000) 'Transferring Knowledge', *Noûs*, 34(1): 131-152.

Greco, J. (1990) 'Internalism and Epistemically Responsible Belief', *Synthese* 85(2): 245-277.

Greco, J. (2004) 'Knowledge as Credit for True Belief', in *Intellectual Virtue: Perspectives from Ethics and Epistemology* (ed. M.DePaul, & L. Zagzebski). Oxford University Press, Oxford.

Green, M. (2007) *Self Expression.* Oxford University Press, Oxford.

Greenough, P. & Pritchard, D. (eds.) (2009) *Williamson on Knowledge.* Oxford University Press, Oxford.

Grice, H.P. (1957) ‘Meaning’, *The Philosophical Review* 66: 377-388.

Grice, H.P. (1969) ‘Utterer's Meaning and Intention Source’, *The Philosophical Review* 78(2): 147-177.

Grice, H.P. (1975) ‘Logic and Conversation’, *Syntax and Semantics* 3: 41-58.

Hanfling, O. (1980) ‘Does Language Need Rules?’, *The Philosophical Quarterly* 30(120): 193-205.

Hindricks, F. (2009) ‘Constitutive Rules, Language, and Ontology’, *Erkenntnis* 71(2): 253-275.

Hume, D. (See Millican).

Hyman, J (2006) ‘Knowledge and Evidence’, *Mind* 115(460): 891-916.

Kelp, C. (2013) ‘A Practical Explication of the Knowledge Rule of Informative Speech Acts’, *European Journal of Philosophy* 21: 367-383.

Kim, J. (1993) *Supervenience and Mind.* Cambridge University Press, Cambridge.

Kim, J. (1999) *Mind in a Physical World.* MIT Press, Massachusetts.

Kolbel, M. (2011) ‘Conversational Score, Assertion and Testimony’, in *Assertion: New Philosophical Essays* (ed. J. Brown & H. Capellen). Oxford University Press, Oxford.

Kvanvig, J (2009) ‘Assertion, Knowledge and Lotteries’, in *Williamson on Knowledge* (ed. P.Greenough & D. Pritchard). Oxford University Press, Oxford.

Lackey, J. (2007) ‘Norms of Assertion’, *Noûs* 41(4): 594–626.

Lackey, J. (2008) *Learning from Words: Testimony as a source of knowledge*. Oxford University Press, Oxford.

Lackey, J. (2011) ‘Assertion and Isolated Second-hand Knowledge’ in *Assertion: New Philosophical Essays* (ed. J. Brown & H. Capellen). Oxford University Press, Oxford.

Leite, A. (2007) ‘How to link assertion and knowledge without going contextualist: A reply to DeRose's "Assertion, Knowledge, and Context"’, *Philosophical Studies* 134(2): 111-129.

Lewis, D. (1996) ‘Elusive knowledge’, *Australasian Journal of Philosophy* 74(4): 549-567.

Littlejohn, C. (2013) ‘The Rusellian Retreat’, *Proceedings of the Aristotelian Society* 113(3): 293-320.

McCulloch, G. (2003) *The Life of the Mind: an essay on phenomenological externalism.* Routledge, London.

McDowell, J. (1980) ‘Meaning, Communication and Knowledge’ in *Philosophical Subjects: Essays on the Work of P. F. Strawson* (ed. Z. van Straaten). Clarendon Press, Oxford.

McDowell, J. (1982) ‘Criteria, Defeasibility and Knowledge’, *Proceedings of the British Academy* 68: 455-479.

McDowell, J. (1994) ‘Knowledge by Hearsay’ in *Knowing from Words,* (ed. B.K. Malital & A. Chakrabarti). Kluwer Academic Publishers, Dordrecht.

McDowell, J. (1995) ‘Knowledge and the Internal’, *Philosophy and Phenomenological Research* 55(4): 877-893.

McDowell, J. (1996) ‘Précis of "Mind and World"’, *Philosophical Issues* 7, Perception: 231-239.

McDowell, J. (1998) *Meaning, Knowledge and Reality*: Harvard University Press, Cambridge, Massachusetts.

Malital, B.K. & Chakrabarti, A. (eds.) (1994) *Knowing from Words.* Kluwer Academic Publishers, Dordrecht.

Miller, A. ‘The Development of Theories of Meaning: from Frege to McDowell and beyond’. Forthcoming in *The Oxford Handbook of the History of Analytic Philosophy* (ed. M. Beaney). Oxford University Press, Oxford.

Millican, P. (2007) *Hume’s Enquiry into Human Understanding*, Oxford University Press, Oxford.

Millikan, R. (2004) *Varieties of Meaning: The 2002 Jean Nicod Lectures.* MIT Press, Massachusetts.

Moran, R. (2005) ‘Getting Told and Being Believed’, *Philosopher’s Imprint* 5(5): 1-29.

Moran, R. (2005) ‘Problems of Sincerity’, *Proceedings of the Aristotelian Society* 105(3): 341-361.

Nagel, J. (2013) ‘Knowledge as a Mental State’, *Oxford Studies in Epistemology* 4: 275-310.

Nagel, T. (1979) *Mortal Questions.* Cambridge University Press, Cambridge.

Owens, D. (2006) ‘Testimony and Assertion’, *Philosophical Studies* 130: 105-129.

Pagin, P. (2008) ‘What is Communicative Success?’, *Canadian Journal of Philosophy* 38(1): 85-115.

Pagin, P. (2011) ‘Information and Assertoric Force’ in *Assertion: New Philosophical Essays* (ed. J. Brown & H. Capellen). Oxford University Press, Oxford.

Pagin, P. (2012) ‘When does communication succeed?’, a PowerPoint presentation, Stockholm University Workshop. URL = http://people.su.se/~ppagin/papers/eurostockholm.

Pagin, P. (2014) ‘Assertion’, in *The Stanford Encyclopedia of Philosophy* (Spring Edition) (ed. E.N. Zalta). URL = http://plato.stanford.edu/archives/spr2014/entries/assertion/.

Pelling, C. (2013) 'Asserting, Telling, and Epistemic Norms', *Australasian Journal of Philosophy*, published online. DOI: 10.1080/00048402.2013.798340.

Pritchard, D. (2004) 'The Epistemology of Testimony', *Philosophical Issues*: 326-348.

Reed, B. (2002) 'How to Think about Fallibilism', *Philosophical Studies* 107(2): 143-157.

Reynolds, S. (2002) 'Testimony, Knowledge, and Epistemic Goals', *Philosophical Studies* 110(2): 139-161.

Ross, A. 1986 'Why Do We Believe What We Are Told?', *Ratio* 28(1): 69-88.

Rowland, M. (2003) *Externalism: Putting Mind and World back Together Again*. Acumen Publishing, Chesham.

Russell, B. (1998: ed. J. Skorupski) *The Problems of Philosophy.* Oxford University Press, Oxford.

Sartre, J.P. (1970: trans. J.P. Fell) 'Intentionality: A fundamental idea of Husserl's Phenomenology', *Journal of the British Society for Phenomenology* 1(2): 4-5.

Schiffer, S. (1972) *Meaning.* Oxford University Press, Oxford.

Schiffer, S. (2001) 'Communication', in *The International Encyclopedia of the Social and Behavioral Sciences* (ed. N.J. Smelser & P. B. Baltes). Elsevier,

Searle, J. (1969) *Speech Acts: An Essay in the Philosophy of Language*. Cambridge University Press, Cambridge.

Searle, J. (1979) *Expression and Meaning: Studies in the Theory of Speech Acts*. Cambridge University Press, Cambridge.

Searle, J., Kiefer, F. & Bierwisch, M. (eds) (1980) *Speech Act Theory and Pragmatics.* D. Reidel, Dordrecht.

Stanley, J. (2005) 'Fallibilism and Concessive Knowledge Attributions', *Analysis* 65:2: 126-131.

Strawson, P. F. (1964) 'Intention and Convention in Speech Acts', *Philosophical Review* 73: 439-60.

Stroud, B. (2000) *Understanding Human Knowledge.* Oxford University Press, Oxford.

Sutton, J. (2007) *Without Justification.* MIT Press, Massachusetts.

Turri, J. (2011) 'The Express Knowledge Account of Assertion', *Australasian Journal of Philosophy* 89(1): 37-45.

Turri, J. (2013) 'Unreliable Knowledge', *Philosophy and Phenomenological Research*, published online. DOI: 10.1111/phpr.12064.

Vendler, Z. (1980) 'Telling the Facts', in *Speech Act Theory and Pragmatics* (ed. J. Searle, F. Kiefer & M. Bierwisch). D.Reidel, Dordrecht.

Weiner, M. (2005) 'Must We Know What We Say?', *Philosophical Review* 114(2): 227-251.

Welbourne, M. (1979) 'The Transmission of Knowledge', *The Philosophical Quarterly* 29(114): 1-9.

Welbourne, M. (1981) 'The Community of Knowledge', *The Philosophical Quarterly* 31(125): 302-314.

Welbourne, M. (2001) *Knowledge.* McGill-Queen's University Press, Montreal.

Williamson, T. (1995) 'Is Knowing a State of Mind?', *Mind*, 104(415): 533-565.

Williamson, T. (2000) 'Scepticism and Evidence', *Philosophy and Phenomenological Research* 60(3): 613-628.

Williamson, T. (2000) *Knowledge and its Limits*. Oxford: Oxford University Press.

Williamson, T. (2004) 'Sosa on Abilities, Concepts and Externalism', in *Ernest Sosa: And his Critics* (ed. J. Greco). Blackwell, Oxford.

Wilson, D. & Sperber, D. (2005) 'Relevance Theory', in *Handbook of Pragmatics* (ed. L. Horn & G. Ward). Blackwell, Oxford.

Wright, C. (1991) 'Scepticism and Dreaming: Imploding the Demon', *Mind* 100(1): 87-116.

Wright, C. (2002) '(Anti-)Sceptics Simple and Subtle: G. E. Moore and John McDowell', *Philosophy and Phenomenological Research* 65(2): 330-348.

Zagzebski, L. (1994) 'The Inescapability of Gettier Problems', *The Philosophical Quarterly* 44(174): 65-73.