The shiftiness dilemma applied to telling and its norms



Candidate Number: 1047537

Word Count: 29998 (excl. Title page and bibliography)

This thesis is submitted in partial fulfilment of the requirements for the Degree of

Bachelor of Philosophy

December 14, 2022

Table of Contents

Abstract……3

Introduction……4

Chapter 1 - Relevant Alternatives to Contextualism to Subject-Sensitive Invariantism

Section 1.1 - Epistemic contextualism and its evolution……8

Section 1.2 - A history of epistemic contextualism……..15

Section 1.3 - The Relevant Alternatives View of Knowledge, The Precursor to Epistemic Contextualism……...16

Section 1.4 - Advocacy for Epistemic Contextualism …….26

Section 1.5 – Epistemic Contextualism Explained……..27

Section 1.6 – Types of Invariantism Explained……...33

Section 1.7 - The Normativity of Assertion……..34

Chapter 2 - The shiftiness dilemma presented……37

Section 2.1 – Introduction……..37

Section 2.2 – How the dilemma emerges……..38

Section 2.3 – How to respond to the dilemma……..41

Chapter 3 - A better way of understanding the contextualist's examples……45

Section 3.1 – The focus of the main argument in this thesis……...45

Section 3.2 – Telling and its norms……...48

Section 3.3 – A more suitable framework for Thelma and Louise……51

Section 3.4 – A better warranted assertability manoeuvre……52

Chapter 4 - Rebuttals of counterarguments……59

Section 4.1 – 4 Strong criticisms and their answers……59

Section 4.2 – Conclusion……84

Abstract

This thesis represents an attempt to argue that there is an invariant epistemic standard that governs the assertability of knowledge attributions, and that any context variability with practical stakes is owed to the interpersonal norm that constitutes the speech act of telling. The way I reach my conclusions is by arguing that the subject matter of epistemic contextualism has been analysed in a sub-optimal way, and that this has led to a significant omission of an important aspect of knowledge attributions and denials, namely their interpersonal normativity. The literature on this issue has been universally analysing all of the examples that have motivated arguments between epistemic contextualists, classical invariantists, and subject-sensitive invariantists as cases of *assertions* of knowledge attributions and denials. I propose that there is a more precise way in which we can analyse the same examples, namely as cases of telling, which will require contextualists and invariantists alike to make room in their analyses for interpersonal normativity.

These two proposals, of analysing examples of knowledge attributions and denials more precisely, and of accepting that telling, as a speech act, involves an interpersonal norm, are not too controversial independently. However, I will use their combination to argue that the context-sensitivity exhibited by uses of knowledge ascriptions and denials owes to the interpersonal norm that constitutes telling, which stipulates that in telling their audience that p, the speaker takes responsibility for the truth of p. I finally argue that the practical stakes involved in a situation determine the amount of epistemic warrant that one needs to have in order to be warranted in taking responsibility for the truth of p.

If all of the arguments presented above are right, I manage to provide a way of evading the context shiftiness dilemma, which is an argument made by the contextualist against the classical invariantist. Moreover, a strong invariantist case has been built.

Introduction

In chapter 1, I offer a detailed explanation of the theory of my main opponent, the epistemic contextualist. By introducing epistemic contextualism[1], along with the contradicting set of intuitions that it purports to give an explanation for, chapter 1 familiarizes the reader in detail with the way in which contextualists propose we should make sense of the expression "know(s)", namely, by treating it as an indexical expression. I then spend a section showing how epistemic contextualism has evolved across the years through an exegesis of landmark papers that have shaped the current state of epistemic contextualism, which I follow with a brief classification of different versions of the view. I finish chapter 1 by introducing the view which contends most fiercely and directly with epistemic contextualism, classical invariantism, explaining the different brands of invariantism. This will set the scene for discussing the issue of the main dilemma that this thesis attempts to explain away and for expanding on the points of contention between contextualists and invariantists.

The straightforward explanation of contextualism I present in chapter 1 will make use of David Kaplan's notions (1989) of content and character. This way of explaining contextualism is suggested by Keith DeRose in his (2008), where he also offers a provisional proposal about the character of "know(s)". While this is not how contextualists themselves ordinarily describe their views, it does well in capturing how epistemic contextualists believe that "know(s) that p" can have different referents in different contexts. I summarize Kaplan's distinction between content and character (in doing so, I also make available to the reader the notions of "circumstance of evaluation" and "context of utterance"), then show how the content and character of the word "know" will, according to the contextualist, work in a way that will

[1] Used interchangeably with "contextualism about knowledge attributions", or just "contextualism" throughout this thesis.

allow for "knows that p" to have variable referents that are determined by context. While ultimately, where one lands on the specifics of the character of "know" is not going to hold much significance for my argument or this thesis, I nonetheless make use DeRose's suggestion of applying Kaplan's distinction to the word "know" as a means of explaining what contextualism about knowledge attributions is about. This will have many advantages, including its generality, the fact that it can be applied to contextualist as well as invariantist theories, that it is easy to understand, and that it does not commit my description of contextualism to any particular contextualist view. I then delve into the particulars of contextualism using the terminology that contextualists like to use, which will be needed for the reader to take on board chapter 2.

Due to the influence it has had in shaping multiple contextualist positions, the significant traction that resulted from it, and the amount of work published on it, I begin by looking at DeRose's well known contextualist position. I take the reader through the well-studied moves by and rebuttals to DeRose's view, which have helped to develop many of the current contextualist views. I also closely examine the ordinary language foundations that DeRose bases his contextualism on, anticipating criticisms I will make against these in chapter 4.

I then move on to the *other* central topic to this thesis, the normativity of assertion. This topic in the philosophy of language constitutes the second half of the dilemma that my argument attempts to refute. Given that the debates between contextualists and invariantists almost entirely turn on one speech act, assertion, I look at the case of assertion in detail, considering the various theories about the normativity of assertion. Assertion is also the speech act that lies at the centre of the shiftiness dilemma, which I spend time detailing in chapter 2, so this chapter will be crucial for understanding the rest of the thesis.

Taking a rigid stance on the normativity of assertion will make it harder for later defending my argument in this thesis, so I will refrain from endorsing any of the suggestions for the normativity of assertion. Even though I remain silent on this issue, no damage will be incurred by my argument as a result. The view that ultimately interests me is the knowledge norm of assertion, as it enables the contextualist to put any invariantist who is inclined to accept the knowledge norm of assertion in check.

Chapter 2 boils down and combines the important information presented in chapter 1 and details the 'shiftiness dilemma' argument made by the contextualist. This dilemma is elegantly used against classical invariantism, which (to jump ahead) is the view that the epistemic standards of knowledge attributions are invariant with practical stakes. The contextualist will argue as such: if invariantists want to argue that the epistemic standards of knowledge attributions are the same in all contexts, then they can't also accept that there is a norm of knowledge governing assertion, given how asserting that p is just as contex-sensitive as knowledge attribution. The contextualist argues that this context-sensitivity that is tracked by their examples proves that classical invariantism cannot be true if the knowledge norm of assertion is true. They will then need to convince the invariantist that their debate can be carried forward from attributions of knowledge like "S knows that p" to simple assertions of "p". If all of the steps in the contextualist's argument are right, the upshot will be defeat for the invariantist. Both horns of the dilemma have a very strong theoretical pull, and the classical invariantist who accepts that assertion is governed by a constitutive norm of knowledge will not be able to either explain the argument away or argue against it, leading to defeat. Chapter 2 explains how the dilemma arises from what appears to be a strong incompatibility between two views: classical invariantism about knowledge attributions, and the knowledge norm of assertion.

While chapter 2 has the purpose of explaining the 'dilemma' argument made by the contextualist, chapter 3 has the purpose of providing an alternative invariantist explanation for the phenomena of shiftiness according to practical stakes. To construct my argument, I first look at the speech act of telling, which will be a central component of my strategy. I introduce this speech act so that I can argue that contextualists and invariantists alike should be analysing all of the imaginary scenarios that motivate the contextualist intuitions primarily as cases of telling. It is telling, not simply assertion, that is the most relevant common denominator to all of the contextualist's thought experiments, I will argue, and their thought experiments should be analysed as cases of telling. There is something important that is lost with this seemingly insignificant difference between telling and asserting. What we lose is crucially the interpersonal dimension of telling, the main difference between the two acts, which is ordinarily excluded in the analysis of the contextualist's examples as cases of assertion. This overlooked interpersonal difference crucially explains, in all of the cases presented by the contextualist, the difference between what a speaker can take responsibility for and what they cannot.

Chapter 4 will be a critical analysis of my argument. I will look at the first possible argument that the contextualist can bring up, and then explain why they do not ultimately refute my argument. I consider criticisms that my view cannot account for why permissibility in high-stakes scenarios can be restored by an increase in one's epistemic standing, that my argument only amounts to a warranted assertability manoeuvre and it is not feasible to claim that permissibility and the epistemic norm of assertion can be completely separate, that I cannot just do what I please with the imaginary scenarios created by contextualists, and that denying the truth of the speakers' claims in the contextualist scenarios amounts to a serious shortcoming.

Chapter 1

-Relevant Alternatives to Contextualism to Subject-Sensitive Invariantism-

Section 1.1 - Epistemic contextualism and its evolution

In this section of chapter 1, I will look at contextualism about knowledge ascriptions as a theory in a general form. Throughout my thesis I will be considering a range of contextualist views about knowledge ascriptions (the views I will be arguing against), so the purpose of this section is to provide a representation of the core of the contextualist thesis. The description of contextualism about knowledge ascriptions that will result at the end of this section will serve as a standard when referring to "contextualism" for the rest of my thesis, as I will frequently have to do so without referring to any one view, and it will only include the common denominator between the contextualist views that my thesis aims to challenge.

The core proposal of contextualism can be neatly drawn out with the help of a distinction introduced by David Kaplan in his (1989) *Demonstratives*, a collection of lectures he gave in the 1970s. Kaplan aimed to develop a semantics for indexicals (terms like "here", "this", or "I", that can change their reference depending on whom they are uttered by, where they are uttered, when, or other factors about the context), and in order to do so, proposed a distinction that proves useful to us, that between the character and the content of an expression. To be able to appreciate the intended difference that Kaplan points to here, it is best to first have a handle on the difference between the notion of the context of utterance of a sentence, and its circumstance of evaluation, which Kaplan makes early on in the same chapter.

What Kaplan calls a "circumstance of evaluation" of a sentence is just a circumstance or state of affairs that one needs to evaluate in order to judge whether a sentence is true, and to judge what or who the terms in the sentence refer to. The context of utterance is simply the circumstance or state of affairs that happens to hold at the time and place of the utterance.

Consider the following two sentences:

1) The Lord Chief Justice is a fair person.

2) Once, the Lord Chief Justice was a fair person.

The circumstances that need to be considered for establishing the truth of (1) differ from those needed for establishing the truth of (2). One needs to establish truths about Lord Ian Duncan Burnett to judge whether (1) is true, and truths about different people, at different times, to judge whether (2) is true. In other words, the contribution that "Once" makes is to change the circumstance of evaluation, so that "Lord Chief Justice" refers to different people in (2), so that some different piece of information needs to be considered for evaluating whether (2) is true. The context of utterance of this sentence, again, is just the situation or state of the world in which the speaker uttered the sentence.

Returning to our purposes of explaining contextualism, we can now make sense of the notions of the content and the character of an expression. The content of an expression is tightly related to its circumstance of evaluation – the content specifies what (or who) that expression refers to, given a circumstance of evaluation. In simpler terms, the content of an expression is just the information that is expressed by a specific utterance[2]. In (1), the content of "the Chief Justice" relevant to its circumstance of evaluation, is the person Ian Duncan Burnett. The content of the same expression "the Chief Justice" relevant to the circumstance

[2] Some writers (Blome-Tillmann, 2022) use the term "referent" to explain the notion of content, as it is simpler, but in terms of Kaplan's semantics for distinguishing between different types of meanings of indexicals, "referent" does not represent its own semantic category.

of sentence (2) will have a different referent, that is unspecified, but excludes Ian Duncan Burnett through the usage of "Once".

As Kaplan goes on to point out, an insufficiency of the notion of "content" is that there are plenty of context-sensitive expressions for which one content will not specify the correct object in all contexts. Expressions such as "here" or "I" are examples of this – these are terms sensitive to features of the context of utterance, and can have different contents even in identical sentences, depending on the place, the time, or the person speaking. This is why Kaplan introduces the notion of the character of an expression. The character is a function that works to determine the content - it is the function mapping the context of utterance to the content of an expression. If I utter "I feel lost" or "Come here", the contents of the terms "I" and "here" are different from the contents of the same sentence when uttered by anyone else (or in any other place, respectively), and it is the character of those terms which determines what the content will be, along with a given context of utterance. An even easier way to grasp the notion is to think of character as what speakers will say about the "meaning" of these context-sensitive terms. For instance, a competent speaker of English will say that the "meaning" of "I" is "the speaker at the context of the utterance" – this is an indication that the "meaning" of any variable character expression, from the speaker's point of view, will just be its character i.e. what they will be able to say that the content is in any context.

Based on these concepts that Kaplan introduces, there are two categories that will be relevant for our discussion about contextualism, into which we can sort expressions. As just noted, context-sensitive expressions have characters that produce different contents depending on the context of utterance, which can therefore be called expressions with a variable character. The expression "here" has a character that can produce different contents, which is determined by the location of the speaker at the time of the utterance. There are also

expressions that will have the same content regardless of the context in which they are uttered, which can be called fixed-character expressions. Such expressions do not always refer to the same object whenever uttered, since a content is also sensitive to the circumstance of evaluation, which can be changed by other operators in the sentence, as we have seen with the addition of circumstance-shifting expressions such as "once", but fixed-character expressions do have the same content in every context.

This brief detour into the philosophy of language has the purpose of offering an easy-to-understand introduction to epistemic contextualism that does not rely on the loaded contextualist terminology. This is now possible, using the concepts defined above: contextualism is the view that "know(s)"is a variable character expression. Expressions of "know(s)" have a character that can result in more than one content, depending on certain factors of the context of utterance. The takeaway here for the contextualist is that "know(s)" can have more than one content, although contextualists are usually a lot more specific about what those contents are. The resulting contents of knowledge attributions[3], according to such advocates, will generally fall along a scale of strength (or a similar dimension) of knowledge, such that the variation in meaning between different uses of "know(s)" is usually a variation in how strongly epistemically positioned the subject of knowledge, S, is to the object of knowledge, p. Since it is hard to be general and still succeed at giving an accurate representation of a complicated view (even harder, a range of views), I will be more specific about how it has been suggested that "know(s)" can act as an indexical expression, which is the classical contextualist view.

One concrete suggestion offered by a contextualist advocate is that the character of "know(s)" could be the following: "S has a true belief that p and [S] is in a good enough epistemic position with respect to p" (DeRose 2008). This provides a broad enough account

[3] Common way of saying "expressions of "knows that p"".

of contextualism, and it allows us to see how different contents can be produced as a result of this formulation. The possibility for variance in the content of "know(s)" results from the formulation "good enough epistemic position": the classical contextualist will colour in this analysis with their own view of what constitutes this good enough epistemic position, and in what contexts, or with their own analysis of knowledge. There are crucially multiple answers to what is a good enough epistemic position with respect to p, on the contextualist view, and this will depend on the context of utterance.

To put more flesh on this proposal for the character of "know(s)", we can imagine a toy contextualist view: for instance, if "know(s)" is uttered in a context where the speaker is wearing a turtleneck, being in a good enough epistemic position with respect to p is to verify that p is true with at least 5 other people i.e. asking them whether p is true, and receiving a positive answer. In another kind of context, when the speaker is not wearing a turtleneck, to be in a good enough epistemic position with respect to p is to have a strong hunch that p is true. There are two possible contents of "know(s)" in this case, which makes this toy semantics a contextualist one. If I say "I know that there are no peanuts in this dish" wearing a turtleneck, the resulting content of my utterance is "I have a true belief that there are no peanuts in this dish and I have verified this with at least 5 other people", whereas I would be saying something else "I have a true belief that there are no peanuts in this dish and I have a strong hunch about it" if I was not wearing a turtleneck. While they will not be as arbitrary, this is the structure of the contextualist views I will be contending against in this thesis. I will deny the contextualist thesis that there are multiple possible contents of "know(s)".

Different kinds of epistemic contextualism vary on several points: they will disagree about which aspect of the context of utterance causes a change in the different resulting contents of "know(s)". Some contextualists will say that the stakes involved (how much practically hinges on whether the knowledge attribution is true) is what determines the

content of “know(s)”. Others will say that it is the environment in which the speaker finds themselves (whether it is a professional/scientific/formal environment, or whether it is more informal and relaxed). This is one point of disagreement.

Contextualists can also crucially disagree on how much variance there is in the possible contents of “knows that p”: according to some views, like the toy one I presented, there are only two possible contents for “know(s)” (these are generally represented as strong and weak meanings of “know(s)”), although these have been generally less successful in getting traction. Other views will say that there is a very wide scope for the content of “know(s)”, due to the wide variation in the possible context of utterance, and that it is difficult or impossible to pin down what “good enough epistemic position with respect to p” amounts to in each situation. However, the variation in the contents of “know(s)” should therefore be seen as a gradient of strength, according to these proponents, which is determined by some aspect of the context of utterance.

On a view such as this one, where the character of “know(s)” contains not a defined set of contents, but rather a gradient where “good enough epistemic position” changes gradually, will run into several difficulties: it must specify where the ‘floor’ or bottom of what a ‘good enough epistemic position is’ can be, such that “know(s)” will have a lower limit on how the sentence in which “know(s)” occurs can still be true for the lowest epistemic position one can find themselves in. Such an advocate would have to say that there is a limit on how a context can affect the content of “know(s)”, such there is no weaker “good enough epistemic position” that can still be a content for “know(s)”. Contextualists will disagree on what that lower limit is, and this will normally be determined by their working theory of knowledge.

A handy feature of the proposal for the character of “know(s)” being “a good enough epistemic position” is that it leaves the possibility open to apply to a wide range of

contextualist views, and even to invariantist views. Nothing about DeRose's suggestion for the character of "knows that p" will prevent an invariantist from claiming that there is an invariant "good enough epistemic position" that is the character of "know". This is therefore a good opportunity to use the same framework to introduce the classical invariantist view about knowledge attributions. Such an advocate will say that DeRose's definition of the character of "knows that p" being "good enough epistemic position" is correct, but they will insist that there is only one way in which one could be in this good enough epistemic position with respect to p, and that this does not change regardless of the context of utterance. On such a view, even if the verbatim definition of the character remains identical to the contextualist's, the character of "knows that p" is not variable, but fixed. This is because of the invariantist thesis that knowledge attributions do not change their content depending on context, so there is only one referent of "know(s)".

Before moving on from this general account of contextualism, it is worth making the following point about taking "knows that p" to be an indexical expression (one with a variable character). In a (2009) paper, Jeff Speaks says that distinguishing between character and content as two aspects of meaning has so far been accepted as uncontroversial in the philosophy of language, as far as it is used for developing a semantics for indexicals. Other applications of this distinction are more controversial, such as using it to explain the difference in cognitive significance between proper names, or other applications that don't immediately pertain to this thesis. However, as we have seen, epistemic contextualism is a thesis that "know(s)" has an indexical semantics, that it is an indexical expression. Nothing controversial has been so far said, as far as using Kaplan's notions for explaining what contextualism amounts to. The only controversial aspect of the thesis that "know(s)" is indexical so far is whether knowledge attributions are, in fact, variable character expressions, or not. However, this is just the form of the invariantist and contextualist debate. In this

thesis, the distinction will only have an informative role – not much hinges on how "knows" fits into the content/character distinction, or what the correct formulation of the character of "know(s)" is, as long as it successfully manages to outline the basic thesis of epistemic contextualism.

Section 1.2 – A history of epistemic contextualism

This section of chapter 1 looks at the beginnings of epistemic contextualism and how it has evolved in its approximately 50 years of being a growing subject in epistemology, to where it is today. I then follow this with a section that looks at the different versions of contextualism, as currently defended, which will be contested by this thesis.

The story of epistemic contextualism is a story of a growing number of epistemologists (i.e. Gail Stine, Alvin Goldman, Stewart Cohen, Keith DeRose) responding to increasing pressures from two directions. The first was from a view of knowledge that was gaining traction in the 1970s and 1980s called the relevant alternatives view of knowledge (Dretske 1970, Sine 1976), with the pressure leading towards the development of epistemic contextualism being the pull to branch off as a natural extension of the view. The second pressure felt was the need for a response to an accumulation of philosophers endorsing classical invariantist and skeptical views of knowledge attributions (Unger 1971), as well as a growing number of arguments in favour of epistemic skepticism generally. As a result, epistemic contextualism was born in the 1970s as a way of answering both pressures: it offered an answer to the then emerging skeptical arguments, provided an alternative to the orthodoxy of the classical invariantist view of knowledge attributions, and it developed as a version of the relevant alternatives view of knowledge. The view started off with relatively few advocates, being more popular and having more proponents and stronger advocation in

the 1990s and early 2000s. In order to get a clear grasp on the chronology of events that lead to the beginning of epistemic contextualism, I first give an account of the relevant alternatives view of knowledge that came out of the 1970s which subsequently branched into a form of epistemic contextualism. I then go on to present one of the first transparently contextualist accounts of knowledge attributions, also touching on the type of skeptical argument that solicited the response of this early contextualist. I end chapter 1 by detailing how the classical invariantist position came into the picture as a thesis about knowledge attributions, as it had already been very much the orthodoxy view in epistemology.

Section 1.3 – The Relevant Alternatives View of Knowledge, The Precursor to Epistemic Contextualism

I will now describe the relevant alternatives view so that it will become easy to understand how early versions of epistemic contextualism sprang from it, initially as somewhat successful attempts at ameliorating the RA view. As mentioned earlier, the RA view is a theory of knowledge, meant to offer an account of the necessary and sufficient conditions for knowing a proposition. It is, of course, through the innovative condition it added to the conditions for knowledge of "true belief", that the RA view set itself apart as a theory of knowledge. Regardless, I will go through all of the conditions put forward by the RA view.

Defenders of RA views will first say that having a true belief that p is a requisite of having knowledge that p. I now look at this first condition in detail. Among those philosophers who take knowledge to be an object for philosophical analysis (there are many who do not), RA theorists will not meet with much opposition on their view of knowledge at this stage, as the true belief condition is still mostly accepted as a sound requirement in

analyses of knowledge in contemporary epistemology. There are exceptions, of course, to this orthodox analysis of knowledge that takes true belief as a necessary condition. There is likewise considerably more resistance to the idea that knowledge the kind of object that should be analysed and instead be considered a basic concept, whereas this view had not gained the kind of traction that we see today.

The RA theory in its early forms adopted the strategy of finding a formula that keeps the "true belief" requirement and adds an additional component (one or more conditions), that seeks to give a full account of knowledge. In spite of seeming like an obviously necessary condition, the truth condition has faced its share of opponents. In his (2010), Allan Hazlett is opposed to the truth condition of knowledge, on the basis of seemingly felicitous uses of 'know' that are not factive (by "factive" Hazlett means uses of "S knows that p", where "S knows that p" entails that p). It is not difficult to come up with instances where this arises in ordinary language: "Everyone knew that smoking was harmless, until…" is one example of such a non-factive use of "know". Hazlett uses this to argue against the factivity of "know", along with various other verbs. He is in a serious minority here; the truth condition is otherwise overwhelmingly accepted as a necessary condition for knowledge by epistemologists.

As concerns the belief condition, it is likewise also largely accepted that one must believe that p order to be able to know that p. Trying to imagine a case of a subject who knows something but who does not believe that same thing will, at first glance, strike probably anyone as at least very difficult. Even so, there are several arguments that have been made against belief as a requirement of knowledge – such as Radford (1966) or Myers-Schulz and Schwitzgebel (2013), who argue that it is possible to construct imaginary scenarios of knowledge in the absence of belief, and that these are not impossible to come up with, and that counterexamples of this type, which they claim to be coherent and possible,

should inform any analysis of knowledge that takes counterexamples as evidence. Still, the orthodoxy in contemporary analyses of knowledge, as the RA view does, accepts the truth condition and the belief condition as uncontroversial. I elaborated on this first condition of the RA view in order to show that while the true belief condition is certainly not blindly accepted or simply taken for granted, it is certainly still the standard, and is present in the RA view as well. I now move on to what distinguishes the RA view, the meat of the theory.

The contribution of the RA theory to epistemology in the 1970s unsurprisingly rests on the addition to this generally accepted requirement for knowing that p: the condition that one's true belief at a time must be supplemented by an ability to eliminate all relevant alternatives to P, at that time. This is how Fred Dretske chose to forumulate his first RA theory, in his (1970), and it is the formulation directly tied to the first developed forms of epistemic contextualism. Several versions of the RA view have been articulated with the points of contention between existing views being about which alternatives should be considered relevant and why, and what determines this variance.

However, what was more of a point of controversy with the above view was Dretske's own interpretation of which possibilities should be counted as relevant, and much of the ensuing discussions of his view concern just this. Dretske believes that we should only include those alternatives that *might* (his preferred formulation) have obtained if the actual state of affairs did not occur. There is a lot to unpack here, and since most of it will not have direct links to our discussion of contextualism, we will put it aside, but it is this issue of determining which alternatives should be included into the notion of relevance that initially motivated the first accounts of epistemic contextualism.

This brings us to the first epistemic contextualist theory. The philosopher who was first to explicitly weave the RA theory with knowledge attributions and turn against the invariantist tide was Gail Stine, who in her (1976) Skepticism, Relevant Alternatives, and

Deductive Closure article put forward a view that, save for much of the terminology that we have now come to expect of contextualists, is one of the earliest clear-cut cases of an epistemic contextualist account. In her paper, Stine was primarily concerned with the relevant alternatives view of knowledge, intending to show how the RA view that she attributed to Fred Dretske could be used to argue against epistemic skepticism. More specifically, she intended to dissolve skeptical hypotheses that aim to question the most ordinary cases of knowledge and ultimately conclude that any such knowledge is unobtainable (though Stine was interested in a specific type of skeptical argument). In doing so, she formed a fleshed-out version of contextualism about knowledge attributions. She merged this formulation of contextualism with her RA view of knowledge in order to construe the skeptic's argument as an argument with no real weight.

The type of skeptical argument that Stine aims to disprove goes along these lines:

P1) Skeptical scenarios to the effect that one is deceived by an evil demon, a brain in a vat, or inside a simulation machine, cannot be completely ruled out.

P2) The above skeptical scenarios or possibilities are not compatible with p.

P3) To be able to know that p, one should be able to rule out all of the skeptical possibilities like the ones mentioned above.

C1) One does not know that p. (Stine 1976)

This type of argument was seen as particularly salient at the time because it stems from the principle of epistemic closure, which motivated the development of the relevant alternatives view. The closure principle says that if one knows that p, and one knows that p entails not-q, then one must also know that not-q. Dretske believed that the RA view was incompatible with closure (1970), whereas Stine aimed to conclude that the RA theorist need not abandon the principle of epistemic closure.

Part of Stine's principle argument against the skeptic was just a form of epistemic contextualism – she held that in almost all circumstances, the alternatives that the skeptic insists we take into account (like the examples mentioned above), have no relevance whatsoever. These are possibilities that only the skeptic will raise, which are constructed in such a way that they can be always brought up, and they should be resisted if raised as relevant, or else not be seriously considered. The possibility that an evil demon may be deceiving us about something we mistakenly (according to the skeptic), take ourselves to know, is one example of such an immaterial skeptical possibility that can be always be raised.

Stine, as epistemic contextualists do, believed that attributions of knowledge have "tighter criteria" for applicability in those contexts where more alternatives become salient. Stine developed her RA view of knowledge by primarily concerning herself with the concept of knowledge itself being amenable to more stringent and less stringent contexts, such that, according to Stine, it is a matter of context what should be counted as 'knowledge'. Her account of the attributions of 'knowledge' simply follows from this. Following her reasoning, only when the skeptic is successful in installing their stringent epistemic standards by making skeptical alternatives salient, will they be also successful in saying that, in those specific contexts, ordinary 'knowledge' ascriptions (e.g. I know that I have hands) are false. These would be true and uncontroversial statements ordinarily, regardless of the context of the utterance. This is to say that Stine concedes that in contexts such as an epistemology seminar where the possibility of being a brain in a vat is brought to salience and this skeptical scenario becomes a relevant alternative, any 'knowledge' ascription which would otherwise be uncontroversial will no longer be true, due to the skeptical possibility made salient. However, Stine does not believe that it is as easy as bringing up the possibility, for the skeptic to be able to install these standards; she thinks that it is much more difficult than the

skeptic would like, and that cases where very stringent skeptical possibilities are actually accepted are very rare.

This may initially be seen as a kind of victory for the skeptic, considering that they manage to get away with saying that there are scenarios (e.g. the philosophy seminar described above) in which statements like "I know that I have hands" are no longer true (as it is in non-skeptical scenarios) which can apply to any ordinarily uncontroversial objects of knowledge, which Stine will accept. Still, she will deny that there is any reason for the skeptic to start celebrating. Even if the skeptic is granted that in scenarios such as these participants lose the ability to claim 'knowledge' over previously uncontroversial objects of knowledge, the skeptic, Stine claims, does not succeed in making any substantive claim about sloppy language, since the mechanism of installing the skeptic's standards is "perverse", in Stine's view (1976). More importantly, however, the skeptic does not manage to use this small victory (of having a limited number of cases in which their denial of seemingly straightforward cases of knowledge are technically true) to make any broader point about any normal cases of knowledge attributions, since this is achieved merely through a contorted instalment of skeptical possibilities. While the aim of her paper was not one of direct advocacy for contextualism about knowledge attributions, she commits herself explicitly to contextualist positions throughout her paper, most plainly exemplified by the following quote: "It is an essential characteristic of our concept of knowledge that tighter criteria are appropriate in different contexts (11). It is one thing in a street encounter, another in a classroom, another in a law court – and who is to say it cannot be another in a philosophical discussion?" (Stine, 254, 1976)

This is a good moment to make the following important point: while epistemic contextualism is a type of theory that makes substantive claims about the context-sensitivity of knowledge attributions, it leaves the matter open on the particulars of the theory of

knowledge that one adopts. I say that contextualism is a "type" of theory because, as we have seen, there are many different versions of epistemic contextualism around, none of which may act as an accurate representation of all the others. In this way, epistemic contextualism can be used with any number of different theories of knowledge, as it says nothing about the nature of knowledge itself. Similarly, as Jason Stanley points out in his (2005), contextualism, as such, also says nothing about the way in which knowledge attributions change their reference from one context to another, or exactly which factors about context may change the reference of knowledge attributions in the way that contextualists claim they do. So it is important not to mistake epistemic contextualism, as such, for a theory that, on its own, tries to say anything about the nature of knowledge.

Following Stine, similarly minded RA theorists were starting to weave the RA theory into epistemic contextualism. Alvin Goldman took a close approach to Stine's with his RA view (Goldman 1975), arguing that to be able to know that p, one must first be able to discriminate between cases in which P is true and cases in which P is false. Goldman primarily had perceptual knowledge in mind here, which further relied on his views about perceptual equivalence between two possible scenarios. For our purposes, the caveat is that only some of the possibilities of P being false are relevant. This means that, according to Goldman, for one's true belief that p to count as knowledge, one must discriminate between cases in which p is true and only relevant cases in which p is false, giving his own account of the factors that can affect the range of alternatives which are relevant. (Goldman 1976). Through his distinctions of factors that can affect which alternatives can be considered relevant, Goldman inadvertently showed how the RA view can be adapted to contextualism as well as to classical invariantist views, while he favoured the view that can be more easily adapted to contextualism. This will turn out to be relevant in discussions concerning whether there is an "attributor" vs "subject" contextualism distinction, which differentiates between

whether contextualist views can disagree on the issue of whose context should determine the change in the truth-conditions of knowledge ascriptions. I will go through these factors and how they relate to epistemic contextualism, taking the side that only "attributor" contextualism should count as "contextualism", but first show where the two views overlap.

It is important to keep in mind that the discussion of the intersection of RA views and epistemic contextualism pertains to cases of knowledge *attributions* (i.e. cases in which a speaker attributes 'knowledge' to a subject), rather than just cases of knowledge. RA theories make claims about when an instance of a true belief is a case of knowledge, whereas epistemic contextualism makes claims about the truth-values of knowledge attributions. Their intersection is therefore the subject matter of epistemic contextualism, i.e. where there are claims made about subjects having knowledge of p. In his exposition of the intersection of RA with contextualism (DeRose 2009), DeRose presents his make-do account of the RA account of knowledge attributions as the following:

"A's claim that 'S knows that p' is true if and only if S has a true belief that p, and S can rule out all the relevant alternatives to p."

This will do for the short discussion below. The reader should bear in mind that this RA account of knowledge attributions pertains to cases in which a speaker ascribes or denies knowledge to a subject, making no assumptions about whether the speaker and subject coincide. It is also important to bear in mind that, while the formulation of this RA view has a contextualist *feel* to it, since it speaks of 'claims of "knows that p"' and the contribution of "relevant" in the formulation may be most easily interpreted as *relevant to the context* (which would then entail that the speaker's context may affect the range of what is relevant to rule out, which in turn entails epistemic contextualism), this still does not make it a contextualist

or invariantist one necessarily, but remains open to both. The invariantist will merely have to maintain that regardless of the context of the speaker, the alternatives which will be needed to be ruled out (i.e. relevant to the subject) in establishing the truth of a claim of 'S knows that p" are always the same. The contextualist will say that indeed, what counts as a relevant alternative that needs to be ruled out in order to establish the truth of "knows that p" can vary from context to context. (DeRose 2009). But we are jumping the gun: let us first list the different types of factors recounted by Goldman, so we can see their importance to epistemic contextualism.

The first class of possible factors that Goldman recounts as potentially able to affect which alternatives should be taken into account towards determining whether a subject, S knows that p, are factors about the context in which the knower finds themselves. If these are the only factors influencing factors to alternative relevance, the details of the situation in which the knower finds themselves (e.g. the time, the place) determines all of the alternatives which are relevant to whether the subject knows that p. This includes any factors about the context in which the knower finds herself.

The contrasting class of factors that can influence the range of relevant alternatives have to do with the attributor of knowledge, rather than the person to whom knowledge is attributed. This only pertains to the context of utterance of the knowledge attribution or denial, even if the attributor is not identical to the person attributed knowledge. Goldman claims that in this class of factors, it is important that we look at the speaker's linguistic and psychological factors.

This detour into Goldman's paper, indicated by DeRose in his (2009), has a two-fold purpose. First, we are able to see how the RA view can be adapted to both invariantism as well as contextualism: if we take the stance that only factors concerning the subject should be counted as able to influence which relevant alternatives to rule out, the resulting RA position

is an invariantist one. To see why, we can notice that the RA position will be just one that equates the truth conditions of "A's claim that "S knows that P""[4], with the truth-conditions of "S knows that P", while denying that factors about the context of A (where A claims that 'S knows that P') have any effect on the truth of A's claim. This just means that the truth of A's claims is established by whether it is true that S knows. There is no change in the truth-conditions of A's claim, regardless of the context in which A finds herself, which exactly what would be needed for the view to be a contextualist. This RA view of knowledge attributions therefore, given that has the same truth-conditions under every context it is uttered in, is invariant.

If, conversely, an RA theorist takes the stance that allows factors to do with the attributor to be counted as able to influence which alternatives to rule out, the resulting position is a contextualist one. For instance, let's say that the speaker (A) claims that "S knows that P" in an epistemology classroom. In this epistemology seminar, the teacher and students (which includes A) are talking about S, an imagined potential subject of 'knowledge', who exists in an imaginary world. In this world, it is universally accepted that skeptical scenarios to the effect that one may not be able to rely on anything one knows, including one's memory, are considered relevant possibilities to be ruled out in general. This is because the people in this world consider it possible for an evil demon to be deceiving them about virtually anything, so this possibility, for them, holds relevance. In a scenario such as this one, the analysis of an RA theorist will be that A's claim that "S knows that P" is false, because of the relevant possibility of the demon in S's world. This amounts to a contextualist RA position, since the truth-conditions of "S knows that P" are not only dependent on factors about the context of S, but also on the speaker of the claim "S knows that P", which can change its truth-value depending on the context of the speaker. As it

[4] where A is not S

happens, Goldman himself opts for the contextualist reading of his relevant alternatives theory, and does not endorse RA views that only take into account subject factors to determine which relevant alternatives must be ruled out for establishing the truth of the claim "S knows that p", and claims that this is the more attractive view.

Section 1.4 – Advocacy for Epistemic Contextualism

If Stine, Goldman and Dretske contributed through discussion of epistemic contextualism as a theory of knowledge attributions through its ties to the relevant alternatives view of knowledge, it was Keith DeRose and Peter Unger who took the view and ran with it as its own epistemological issue and cemented future debates as important in their own right. Unger's positions changed considerably several times, before somewhat settling in 1990 to the account that we can recognize as his own.

DeRose references Peter Unger as one of the first epistemic contextualists, also crediting him with coming up with the term "invariantism" (2009). It is skeptical views about one's ability to know anything that the RA theorists were contending against when developing their form of epistemic contextualism, the latter being a reaction, which ultimately turned into its own theory about knowledge attributions.

Peter Unger has changed his view on contextualist matters more than twice, which is why DeRose categorizes the history of his changes of heart as *early*, *middle*, and *late* Unger. For our purposes, here is the short of it: early Unger was a skeptical invariantist (1975). He claimed that the standard for the truth of knowledge ascriptions is invariant – nothing about the speaker's context can change the meaning of "knows". Still, this invariant standard that must be met for truth claims to be true, per early Unger, is too high to ever be met. The upshot is that most knowledge ascriptions turn out to be flatly false. What early Unger

concedes is that there is variability in when one is *warranted* to ascribe knowledge to a subject. Truth-conditions and conditions of when it is warranted to claim "S knows that P" come apart, per Unger.

Middle Unger (1984) is the period in which Unger abandons his hard-line skeptical view to endorse a more relativistic view, and where DeRose finds the first use of the word "invariantism" to mean "invariantism about knowledge attributions". This relativism of Unger's is about the meaning of "know(s)" – he claimed that there is no matter of fact over which is truer of his earlier skeptical invariantism or contextualism about "know(s)". This position tows a difficult line, but in spite of this, this period of Unger's trajectory saw the most complete outline of contextualism up to that point in time.

A year later, DeRose distinguishes a Late Unger period, in which Unger finally jumps ship to the contextualists, (1986), abandoning his earlier invariantist, relativist or skeptical view points about knowledge attributions. It is this at this point that we are able to see epistemic contextualism develop into the mainstream, as writers like Steward Cohen started writing about contextualism (Cohen 1986, 1987). It is with DeRose's (1992), Cohen's work, and Lewis's seminal (1996) when contextualism really picked up, and exponentially grew with Williamson's (1996) and (2000) works defending the knowledge norm of assertion.

Section 1.5 – Epistemic Contextualism Explained

At last, we reach the point where I can explain, in more specific terms than through differences in the *content* and the *character* of the term "know(s)", in what consist the different varieties of epistemic contextualism. Since he offers such useful examples and champions his view fiercely, I start with Keith DeRose's. DeRose made a mark with his 1992 paper *Contextualism and Knowledge Attributions,* in which he introduced a well-known

scenario he calls the *Bank Cases*, which have helped propel epistemic contextualism into mainstream philosophy. These have been since used to explain his contextualist position, or contextualism generally. As they stand, however, the Bank cases do not do DeRose's current view justice, due to the relative simplicity of the example which does not properly discern between the attributor and subject of knowledge.

A more helpful and up to date example offered by DeRose is one he offers in his (2009) book, which he refers to as *Thelma and Louise.* This example is well-constructed and long, but this allows us to pinpoint exactly what epistemic contextualism claims about the different contributions that the word "know(s)" can have in different contexts. It is therefore a very helpful way to show the reader exactly the way in which factors about the context of the speaker can influence the epistemic standards that govern usage of the word "know". Here is the example:

Case Set-Up: The Office. Thelma, Louise, and Lena are friends who all work in the same office. Today is their day off, but, before getting an early dinner together, they decide to walk up to the office to pick up their paychecks. Thelma and Lena are also interested in finding out whether a certain colleague is at work, as they are involved in a small office bet with some other workers over whether the often-absent John would show up today. As they pass the door to John's personal office, they see his hat hanging on the hook in hallway, which, in their long experience, has been a sure-fire sign that John is in fact at work. They also hear one working colleague shout to another, 'Why don't you clear that letter with John quick before you send it off?' Satisfied that John is at work and that Thelma and Lena, who bet that he would be, are in a position to collect their winnings from some other office workers, the three friends pick up their checks, go out to dinner together, and then part

company, Thelma going to a local tavern to meet other friends, and Louise and Lena each heading in different directions to go home.

Thelma at the Tavern. At the tavern, which is renowned for the low epistemic standards that govern the conversations that take place within its walls, Thelma meets a friend who bet that John wouldn't be at work, and so owes Thelma $2. Thelma says, 'Hey, John was at work today. Pay up, sucker!' When her friend asks, 'How do you know?,' Thelma replies, 'I went up to pick up my paycheck this afternoon. His hat was hanging in the hall outside his office door, and I heard Frank telling someone to quickly check something with John before sending it off.' Satisfied with Thelma's evidence, the friend pays up. Then, wondering whether Lena will know to collect what she is owed by yet another worker, Thelma's friend asks, 'Does Lena know that John was in?' Thelma answers, 'Yes, she was with me. She knows, too.' Meanwhile . . .

Louise with the Police. Louise has been stopped by the police on her way home. They are conducting an extremely important investigation of some horrible crime, and, in connection with that, are seeking to determine whether John was at work that day. It emerges that they have some reason to think that John was at work and no reason for doubting that (other than the fact that he is often absent from work, which Louise already knows), but as the matter has become so important to the case, they are seeking to verify that he was there. When the police ask her whether she could testify that John was at work, Louise replies, 'Well, no, I never saw him. I could testify that I saw his hat hanging in the hall, which is a very reliable sign that he's at work. And I heard Frank Mercer telling someone to check something with John, as if John were in. But I suppose John could have left his hat on the hook when he went home some previous day. And though it would be a bit strange for Frank not to know whether John was at work, especially that late in the day, I guess he could have been just assuming John was there because John was scheduled to work—and because his

hat was in the hall. You should check with Frank. He at least seemed to know that John was in.' When the police ask Louise whether Lena might know that John was in, Louise replies, 'No. She was only at the office very briefly, with me, and didn't see John, either. She has the same reasons I have for thinking John was there, but, like me, she doesn't know that John was there.'

Crucially, the contextualist position adds the following two claims, endorses them, and seeks to explain why they take them to be true:

When Thelma says at the tavern that Lena knows that John was at work, her claim is true.

When Louise tells the police that Lena doesn't know that John was at work, her claim is true.

(DeRose 2009)

I will now go through this example thoroughly, with the purpose of explaining DeRose's contextualist position. Doing so will allow me to isolate the contribution that "know" makes in the two use-cases, which will reveal how the contextualist takes the change in reference between the two cases to occur.

Several things should be noted about this example, which have been deliberately picked by DeRose to best illustrate his contextualist position. Firstly, we have our subject and our attributor(s) separated. Lena is the person to whom knowledge is attributed and denied, respectively, while Thelma and Louise are the only two people who attribute and deny knowledge. This is a feature that was not present in DeRose's original bank cases (1992). Secondly, Thelma and Louise are attributing and denying knowledge to the exact same

person (Lena), with all of the details of their knowledge ascription being exactly the same, including the amount of epistemic warrant all of the protagonists have.

What this example very neatly shows is that we can very easily devise examples in which the same subject and proposition can be referred to at the same time, while one speaker can truthfully attribute knowledge to them, and the other, can simultaneously deny knowledge of them. Since the only thing that differs are the practical stakes in Thelma's case versus Louise's case, we can conclude, per the contextualist, that the contribution of "know(s)" is such that it can change its reference depending on the context of the speaker, which is just what epistemic contextualism amounts to: the view that the truth-conditions of knowledge attributing and knowledge denying claims are sensitive to certain aspects of the context of utterance. We have very good evidence that the reference of "know", is in fact sensitive to practical matters.

In Thelma's case, the truth-conditions of a knowledge ascribing assertion are determined by the lower epistemic standard, which is determined in turn by her context. These are different from the truth-conditions that apply to knowledge denying claims of Louise since a higher epistemic standard is set by her context. Their different contexts make it such that the same subject can be ascribed and denied the same predicate at the same time: "know" in the sentence uttered by Louise means something different from what "know" means in Thelma's utterance, the contextualist will want to say.

There are, of course, different brands of contextualism, with varying degrees of uptake. David Lewis, for instance, holds a similar contextualist position, in as much as he is also committed to the position that the truth-conditions of knowledge-ascribing claims can change according to context. His explanation stands in contrast to the one proposed by DeRose: for Lewis, what determines the changes of the truth-conditions of knowledge-ascribing sentences according to context are just changes between which possible worlds are

included in the set that needs to be eliminated in the different contexts, in order for the subject to know, according to Lewis, that P is true (Lewis 1996). Since "knowing that P is true" is for Lewis a matter of P holding in every possible world left uneliminated by the evidence the subject has, the way that the reference of "know" can change is if the extension of "every" in "every possible world" can change in different contexts. This is importantly a distinct approach to DeRose's since DeRose is more concerned with the propriety of knowledge-ascribing claims, rather than with the nature of knowledge itself, on which DeRose does not spend as much energy. For Lewis, on the other hand, his brand of contextualism is a consequence of his broader views of knowledge.

Other versions of contextualism take more concrete epistemological stances in comparison to DeRose's. Robert Hambourger claims that what determines the shift in the truth-conditions of knowledge ascribing sentences is a change in the amount of evidence needed to know a proposition (Hambourger, 1987). Here again, the amount of evidence needed for knowledge is context-sensitive, so the shift in the truth-conditions of knowledge-ascribing sentences is a product of a change in what counts as knowledge, making it a distinct position from DeRose's.

Not all contextualists will agree on the set of factors that determine the variability in the reference of "know(s)". However, for ease of argument and not having to make multiple caveats to each step of my argument in chapter 3, and due to the fact that I concern myself primarily with more fleshed out views, such as DeRose's, I will simply refer to the *practical stakes* involved in a scenario (how much or how little hinges on p) as what determines the variability in the reference of "know(s)". I acknowledge that this does not reflect the whole spectrum of contextualist views, but it will have to do for my argument in chapter 3. In the interest of space, we now move on to the rival of contextualism.

Section 1.6 – Types of Invariantism Explained

Invariantism is the direct competitor of epistemic contextualism, and it states that the epistemic standard of knowledge attributions and denials is invariant. Nothing about the context of the situation can cause any change in what it takes for a knowledge claim to be true. There are other ways that invariantism has been described, including that the truth of knowledge claims is insensitive to practical concerns (Mona 2021). There are multiple types of invariantism, discussed below.

Classical invariantism (sometimes referred to as 'insensitive invariantism') is just the view described above. As such, classical invariantism makes no attempt to vindicate or account for the intuition that speakers in the positions of Thelma and Louise seem to be somehow right in their knowledge ascription and denial, respectively, which seem to be determined by aspects of the context. It just claims that one of the two speakers have to be speaking falsely. That being said, the usual way in which someone who wishes to adopt classical invariantism can argue against epistemic contextualism and account for the context-sensitivity presented in the contextulist examples is through a warranted assertability maneuver: arguing that what the contextualist is right about is pointing out that the use of knowledge attributions is a context-sensitive matter. However, what changes with the high stakes and low stakes cases is the standard under which it is *appropriate* to claim to know. DeRose's example, on such a view, tracks the propriety of knowledge attributions and denials, not their truth-conditions or truth values.

Moreover, as Simion points out (2021), classical invariantism is an orthodox epistemological view. It fits in neatly with other fields of epistemology, and according to Fantl and McGrath, (2009) "anyone who has even passing knowledge of analytic

epistemology in the last fifty or so years knows that moderate invariantism is the orthodoxy. It is the view to beat'" (Fantl and McGrath 2009).

Subject-sensitive invariantism is a form of invariantism that seeks to account for the intuition that more epistemic warrant is required for the truth of knowledge attributions in high-stakes scenarios, and that less epistemic warrant is required for the truth of knowledge attributions in lower stakes scenarios. Jason Stanley in his *Knowledge and Practical Interests* (2005), puts forward a view that stipulates that the reference of "know(s)" does indeed remain invariant, but what can change according to context is what it takes for a subject to actually "know" that p, in a given situation. Here, it is the truth-value of knowledge attributions that changes depending to the practical stakes, but only taking into account the context of the *subject*. More specifically, whether a subject knows that p depends on how much that subject cares about the truth of p, or what is at stake for S.

Section 1.7 - The Normativity of Assertion

This section has the purpose of explaining the second half of the shiftiness dilemma, the knowledge norm of assertion. This will require a very brief introduction into the normativity of speech acts – the idea that illocutionary acts like assertion or commanding are subject to norms that regulate their use in specific ways. After the brief overview of speech act normativity, this section swiftly moves its focus specifically to the normativity of assertion, and the arguments that have been made for the different possible constitutive norms of assertion. This is needed in order to introduce KNA - the knowledge norm, which is a central concept for chapter 2 and indeed, the rest of my thesis. The knowledge norm of assertion is a thesis about the constitutive norm of assertion – it states that one ought to assert that P if and only if one knows that P. This section has the purpose of making sense of this

proposal of a constitutive norm of assertion, and presents the various other views about the norm(s) of assertion, showing that the knowledge norm is controversial, which will contribute to making the contextualist's shiftiness dilemma vulnerable to criticisms to the knowledge norm.

One of the lasting contributions of 20th century analytic philosophy has been, from multiple directions, the advocacy of the idea that speech ought to be seen as more than an activity for sharing information or representing the world, but primarily as something that speakers engage in to perform a broad range of actions: e.g. to wed, to promise, to warn, to condemn. Tied to this new way of looking at speech was the idea that the performance of these actions is subject to rules and norms that govern their use (Austin 1962, Searle 1969). It is not difficult to entertain the notion that some kinds of rules or norms are part of speech acts when one considers ordinary speech practices: e.g. speakers are held to account if they make assertions which they do not believe, which indicates that belief is in some way involved in the norm governing assertion. We expect that speakers who ask questions do not know what the answer to the question is, which suggests that a lack of knowledge of p involved in the normativity of questioning. For instance, it would constitute a breach of a norm involved in asking if I was with a friend, with an apple in front of us, which both of us could clearly see, and I asked "Is there an apple in front of us?", which shows something about the lack of knowledge about p that is involved in a speaker S asking a question about p. Normativity seeks to explain how this happens, and how to understand these rules and how rigorously or not they are followed.

Before moving on to making sense of these norms, it is important to point out that "speech act normativity" in the way that I use it, refers only to the performance of speech acts. When it comes to the rules of syntax or semantics, we are able to speak of rules as well, but we are concerned with a different kind of normativity: the normativity of speech acts.

Briefly, while the literature is still evolving, there is some consensus on the types of rules that regulate speech acts: Searle (1969) stipulates there are constitutive rules that govern speech acts, which constitute the activity which is dependent on them, without which the activity itself could not exist. He also stipulates that there are regulative rules, which specify and regulative moves within an activity, more specifically how the speech act should be performed.

There is still fierce ongoing debate concerning the constitutive norm of assertion, with the topic being very much alive in the philosophy of language and epistemology. Still, Williamson's (1999) and (2000) have left a mark when it comes to debates between epistemic contextualists and invariantists, as he introduced his Knowledge Norm of Assertion (KNA), which states that one can only permissibly assert that p if one knows that p. There have been plenty of other proposed constitutive norms of assertion, including the norm that one should only assert that p if p is true, the truth norm of assertion. This has recently been defended by MacFarlane (2014) and Whiting (2012). There is also the justification norm of assertion, which says that a subject should only assert that p if they rationally believe that p. (Douven 2006) (Gerken 2017). Williamson aimed to identify the constitutive norm of assertion in his (1996) by considering potential norms and finding failures in them, including the truth norm of assertion, and the warrant rule of assertion, until reaching the knowledge norm of assertion. Given the level of debate in the field, the knowledge norm, in spite of its traction, is still a controversial view. In the next chapter, I put together the information presented in the sub-sections of chapter 1, and present the context shiftiness dilemma.

Chapter 2

-The shiftiness dilemma presented-

Section 2.1 - Introduction

In the previous chapters, I introduced the reader to my opponent's theory, contextualism about knowledge attributions, along with the history that has shaped it, and the versions to be found today. I focused on a general contextualist theory, presenting DeRose's example of Thelma and Louise. I went on to present invariantism, the opposing viewpoint concerning knowledge attributions, which will result as a consequence of my argument later in the thesis.

With these notions in our arsenal, we are finally equipped to get to what has been named "the shiftiness dilemma", which will be the focal point of this thesis henceforth. This dilemma was initially presented by DeRose as a strong counter-argument against classical invariantism (DeRose 2002). This chapter tells what this dilemma is all about.

"The shiftiness dilemma" has been coined as such to indicate that it is a dilemma that aims to account for the shiftiness intuition: the intuition that in higher stakes scenarios, more epistemic warrant is needed for permissible assertions to be made, and that in lower stakes scenarios, less epistemic warrant is required for permissible assertions. In her *Shifty Speech and Independent Thought,* Simion (2021) presents a systematic way of understanding the shiftiness dilemma, coming up with short abbreviations for terms that occur often in the debate. Since they prove to be a rather efficient and short way of understanding the crux of the debate, I will make use of them throughout this chapter.

One of these abbreviations coined by Simion, I have already used, which is the Shiftiness Intuition. This lays the groundwork for the strong pull of DeRose's contextulaism,

and will generally be accepted by most people, since it just stipulates that there is an intuitive pull of DeRose's examples.

Section 2.2 – How the dilemma emerges

The first step to understanding how the dilemma emerges is to remember both classical invariantism and the knowledge norm of assertion. Classical invariantism states that there is an invariant epistemic standard that determines the truth of knowledge claims. The knowledge norm of assertion states that one should assert that p iff one knows that p (Williamson 2000).

We have seen from the *Thelma and Louise* case that, indeed, the shiftiness intuition aligns with ordinary speaker's intuitions about how permissible assertions shift according to practical stakes. This happens not only when it comes to knowledge attributions, but it happens for nearly all assertions. In the Thelma and Louise cases, not only is it impermissible for Louise to say "Lena knows that John was at the office", it is equally impermissible for her to say "John was at the office", since she has the same amount of epistemic warrant as Louise does. Likewise, it is permissible for Thelma to say "Lena knows that John was at the office", just as it is for her to assert "John was at the office". So this is not something unique about the permissibility of knowledge attributions and denials. This gives classical invariantists very good reason to protest against the contextualist: their explanation of how the reference of the term "know(s)" can shift reference depending on the context cannot for the larger phenomenon that we can observe in ordinary language.

Contextualists have been able to come with a very good response to this Generality Objection, by using Williamson's knowledge norm of assertion. They will argue that, if there is a constitutive norm of assertion such that one's assertion is only proper if one knows that p,

then surely this means that the propriety of one's assertion that p will always be determined by one's epistemic standards for knowing that p. This offers the contextualist an easy opportunity to generalize their point about the context-sensitivity of knowledge attributions and denials to all assertions, by stipulating that behind every permissible assertion that p, there will be a speaker who "knows" that p, and this is what is at the root at the context-variability of assertions.

By combining these two views, the contextualist will be able to cleverly account for the shiftiness in the permissibility of any assertions that p, and not only cases like Thelma and Louise. In her recent (2021), Mona mentions two examples, called Aspirin High and Aspirin Low. Since these will be helpful in demonstrating how the contextualist will be able to use KNA to account for the generalized phenomenon of the shifting permissibility of assertions with practical stakes, I will be using these examples to show the strength found in combining contextualism with KNA:

ASPIRIN-LOW. On the way to your place, your sister lets you know that she has a minor headache and asks you: 'Do you have aspirin at home or should we stop at the pharmacy?' You remember having bought aspirin last month and tell her: 'Don't worry, we have aspirin at home.'

ASPIRIN-HIGH. On the way to your place, your sister lets you know that her two-year-old baby urgently needs an aspirin on doctor's orders and asks you: 'Do you have aspirin at home, or should we stop at the pharmacy?' You remember having bought aspirin last month. Even so, you think it best to have your partner double check and so you call. When you find out that s/he is not at home you say to your sister: 'Well, let's drop by the pharmacy, just in case.' (Simion 2021)

Here, the contextualist will be happy to say that the difference in the practical stakes involved in the two examples makes it the case that there is a lower epistemic standard required for permissibly saying that "I know we have aspirin at home". In turn, the fact that there is a lower standard of permissibility for knowledge claims in Aspirin-Low, this will have a direct impact on the permissibility of simply asserting that "p", since the permissibility of assertion is directly and only linked to the permissibility of knowledge claims. This is what makes it permissible for the subject to outright assert "we have aspirin at home". So, the contextualist can create a link between one's epistemic standing towards p, the epistemic standard for permissibly ascribing knowledge that p, the practical stakes involved, and how warranted one is in asserting that p. In low stakes cases, therefore, one needs a lower amount of warrant to be able to assert that p permissibly.

Similarly, in a high-stakes scenario, such as Aspirin-High, it becomes impermissible to assert that p ("there is aspirin at home") with the same epistemic standing, given that the practical stakes have changed the epistemic standard required for the permissibility of knowledge claims of p, and therefore also assertions of p. There seems to be a 1:1 tracking of when it is permissible to assert "I know that p" and "p", and we indeed see this play out in ordinary language, as more epistemic warrant is required for proper assertions that p.

This is a very good response, and it is able to deal with just about any example you throw at the contextualist. Even more, it seems to put the invariantist in a tough place that they cannot argue themselves out of, which is how we get to the shiftiness dilemma. Here is why this will stump the classical invariantist, according to the contextualist: If we accept the knowledge norm of assertion, and we accept that practical stakes do influence the permissibility of assertions, then we can no longer deny that the epistemic standards for knowledge attributions are invariant. Accepting the shiftiness intuition amounts to simply accepting that the permissibility of assertions changes according to the practical stakes

involved in a context. It then follows that the only property than can cause this shift in the permissibility of assertion has to be one's "knowledge" that p is true. We are then forced to say that by adopting a view where knowledge claims have an invariant epistemic standard, we cannot account for the shifting permissibility of assertions that p. So if we accept that KNA is true and if we accept the shiftiness intuition, we cannot go for a view like classical invariantism. We therefore have to choose: either KNA is true and we have to explain the variability in propriety of assertions that p through something other than classical invariantism, OR we accept classical invariantism and abandon the knowledge norm of assertion, having to account for the shiftiness in the propriety of assertions that p through some other mechanism or norm, other than one's knowledge that p is true. That is the shiftiness dilemma, as put forward by DeRose.

Section 2.3 – How to respond to the dilemma

Where do we go from here? As Simion points out (2021), most of those writing on matters related to the shiftiness dilemma, since this argument was made, have chosen one horn over the other. There are those who go for the classical invariantism horn, who then have to account for the shiftiness in the permissibility of assertion through some other way than the knowledge norm. Usually, such theorists opt to go for a hybrid view that allows for variability to come in through what it takes to "know". They allow for practical matters to influence what can actually count as knowledge, and the practical stakes therefore can influence the subject of knowledge claims. Subject-sensitive invariantism (Stanley 2005), as we have seen in chapter 1, is one such view, but there are others who put forward views that can be called "pragmatic encroachment" views (Fantl & McGrath, 2002), which argue that subject's practical situation determines that standard required for them to actually know that

p is true. The name is meant to indicate that pragmatic considerations (that have to do with contextual factors of the subject) can encroach on the epistemic standard for knowledge. On the opposing side, there are those who accept the knowledge norm of assertion and deny classical invariantism. These theorists will usually argue that the context-sensitivity that leads to variability in what one can permissibility assert will come in through the variation in how one can fulfill the epistemic norm of assertion. The most straightforward explanation, and the view we have spent most time looking at, is just epistemic contextualism. One can comply in multiple ways to the epistemic norm of assertion (the knowledge norm) through the multiple possible references of "know" that are determined, in turn, by the practical stakes involved in a scenario. Another way to go for those who prefer the KNA branch of the shiftiness dilemma is to say that there is a variability in what norms govern assertion, depending on the context (Greenough 2011). Greenough (2011), for instance, puts forward his "Norm-Relativism" view, according to which it is a relative matter, determined by one's conversational context, what norm will govern one's assertion. He argues that in high-stakes scenarios, the norm governing assertion will be different and much more demanding, compared to the norm governing one's assertion in a low-stakes scenario. For instance, very high-stakes scenarios can require that one be able to remove all the possibilities that p is not true, whereas very low-stakes scenarios can only require that one be justified in believing that p is true.

There is however, another way to go. We can find weaknesses in the contextualist's proposed dilemma and argue against it. This is the route that Simion opts for. To begin with, Simion argues that there is no strict incompatibility between KNA and classical invariantism. This is because there are other ways in which an assertion can be proper without meeting its epistemic norm, but meeting other norms. For instance, suppose prudential norms may force one to assert something false if one had a gun to their head, and in that case it would be improper to assert something only because it met the epistemic norm of assertion. This leads

Simion to say that for the dilemma to work, the incompatibility has to be between the knowledge norm as applied to *epistemically* permissible assertion and classical invariantism, rather than simply permissible assertion. As restated, there is indeed a strict incompatibility between the KNA and classical invariantism, but only if the variation of when a speaker is warranted in asserting that p is fully explained by the epistemic norm of assertion (the knowledge norm). As I understand DeRose, this is how he wishes his thesis to be understood, and I do agree that this is what it would take for his proposed incompatibility between KNA and classical invariantism to hold.

However, Simion has serious doubts about the soundness of the idea that the shiftiness intuition shows us the epistemic propriety of assertion. I do too, as I will show in chapter 3. Her doubts stem from the many different norms that can be at play when it comes to any particular assertion, and how these norms can be in conflict – this makes it more difficult than the contextualist would like to isolate which norm drives the intuition. She claims that it is much more likely that what our intuitions "track" in DeRose's bank and Thelma and Louise cases are *all things considered* propriety. Moreover, she gives good reasons against the thought that the shiftiness intuition tracks epistemic propriety. One of these is that she negates what she calls "Content Individuation", the idea that a norm that influences the amount of epistemic warrant required for being warranted in asserting is thereby an epistemic norm. There are good reasons why Content Individuation is not true, and in chapter 3 I will indeed be making my case as one that proposes an interpersonal norm that does affect the amount of epistemic warrant required for permissible telling, which I will argue is not an epistemic norm. The reasons Simion offers against Content Individuation are that a norm affecting a certain property does not thereby make the norm into a norm about that property. Simion puts it better herself here: "when a certain type of action is governed by a norm (e.g. a norm N) that requires a certain degree of a gradable property G and an

overriding norm modifies the threshold for all-things-considered proper action up or down, this does not turn the overriding norm into a norm of the type that it overrides.". I buy into this idea, and this will be indeed a notion that my argument rests on, as many warranted assertability maneuvers will also have to if they will want to make any claim about how the shiftiness intuition tracks overall assertability rather than truth conditions, and want to offer a potential norm that affects assertability rather than the truth-conditions of assertion.

To tie this all together, this is how denying Content Individuation leads to a denial of the shiftiness dilemma, according to Simion: in order to make epistemically proper assertions, according to the knowledge norm, one needs to have knowledge, which requires a certain amount of epistemic warrant. According to the classical invariantist, this amount of warrant is fixed, and according to the contextualist it can vary with stakes. In either case, warrant is gradable, as one can have a range of epistemic warrant. If that is the case, then other norms that are involved in a particular case of an assertion (practical norms, politeness norms, moral norms) can override the epistemic norm and modify the amount that is required for a particular instance of all things considered warranted assertion. As Simion argues, the fact that these other norms, e.g. moral norms, can influence the amount of epistemic warrant one needs to have to assert properly, does not them make them epistemic norms. This means that classical invariantism can coexist with the knowledge norm of assertion. The way this can happen is if various other norms involved in a particular instance of telling can bring the amount of epistemic warrant needed for proper assertion down or up. So, we can stipulate that there is an invariant amount of epistemic warrant that is required for the truth of knowledge claims, and that assertion is governed by a constitutive norm of knowledge. The way in which we can account for differences in permissibility with differences in stakes becomes simple: the variation is accounted for by other norms at play that influence all things considered permissibility.

I agree with this, and do think that the variation in assertability presented in DeRose's examples is not accounted for by epistemic norms. However, this is still not the full picture, and we have no means by which to say, on a case by case basis, what accounts for the higher or lower amount of epistemic warrant for assertability, when this differs to the invariant standard for knowledge. In the next chapter I offer a full account of exactly what accounts for the variation in assertability, in every case, arguing against epistemic contextualism. I provide a classical invariantist account that can explain DeRose's examples, which can coexist, as we have seen, with the knowledge norm of assertion.

Chapter 3

-A better way of understanding the contextualist's examples-

Section 3.1 – The focus of the main argument in this thesis

In this chapter I make my main argument against epistemic contextualism. I focus on a specific overlooked aspect of the contextualist argument which I think has allowed the contextualist's argument to thrive for as long as it has, which is the too vague analysis of the contextualist's example as cases of assertion. I will follow Simion, who says that the classical invariantist need not give up the knowledge norm to have their way, but this does not constitute the strength of my points (even though this is what counts as a win in favour of the invariantist and against the epistemic contextualist). I will say that imaginary scenarios such as Thelma and Louise, the Bank Cases, along with any other thought experiments that seek to validate the truth of contextualism have an overlooked and important thing in common: they do not *just* involve examples of assertions of knowledge attributions or denials, they are examples of speakers *telling* something *to* their interlocutor. The way we should understand

the underlying speech act, therefore, should more closely match what happens in conversation.

In short, the argument presented in this chapter will focus on the main point of difference between telling and asserting. The two speech acts have so much in common that they can be mistaken for each other, which very likely is why telling has been overlooked so decisively as unimportant in this debate until very recently. Telling, as has been argued by so many, is a constative speech act that is unseparated from conversation, which crucially involves at least two interlocutors. This fact about it makes it more complex than assertion: it adds a interpersonal factor or norm that is not present in asserting as such (as one need not necessarily assert *to* someone)[5], and for good reason. When speakers talk to each other, they do more than assert things in the other's direction. They want to assure each other, be assured, be responsible for the truth of what they tell each other, and be accountable to one another. There have been multiple suggestions about what this extra interpersonal factor present in telling but not asserting can be, which include some of the actions listed in the previous sentence. I have not come up with one myself, so I will simply pick my favourite. It is not crucial where one goes from here, however, because regardless of the particulars of this additional interpersonal factor, my argument will stand: whether it is an assuring the other that p, taking responsibility for the truth of p, or making a promise to the hearer that p is true, what explains the shiftiness intuition are the changes that the interpersonal norm can make to the overall epistemic warrant that is needed to make permissible assertions that p. In all of the high-stakes cases presented in the contextualist/invariantist debate, the denial of knowledge owes to speakers wanting to avoid a breach of the interpersonal aspect of telling – we need not reconcile the truth of what each speaker says when they ascribe or deny knowledge, while

[5] To avoid confusion: typically, cases of assertion will involve some kind of audience. Nevertheless, my broader point concerns the speech act itself, and the fact that an audience is not required for an instance of assertion, but is necessary in the case of telling.

we also need not say that they don't know what they "mean" or are talking about if they are not saying literal truths. We explain the denial of knowledge in the higher-stakes case through the interpersonal norm of making a promise for the truth of p overriding the epistemic norm, even if the epistemic standards needed for knowledge are strictly speaking, met. Flatly, we will deny that one of the speaker's claims is literally true, but that is not such a bad bullet to bite, when we grant that there is good reason to do so.

This thesis is very far from being the first place one can find arguments against the shiftiness dilemma, or, for that matter, arguments that, as far as the speech act of telling stands, finds that the contextualist's argument is unable to provide an adequate explanation. For this reason, I will show the reader what other arguments that focus on telling have been made against epistemic contextualism, especially one made by Charlie Pelling in his (2014) article. The upshot of these is the same as mine as far as the success of classical invariantism is concerned, but I do not agree with the way that Pelling gets there. Furthermore, the details of the arguments matter, in my opinion, and the differences between other arguments against the shiftiness dilemma and mine are not in any way inconsequential, and should be given attention. This is one of the reasons why it matters how we reach the invariantist conclusion.

We should not dismiss or lose sight of this additional component that telling brings into the picture. It may make our understanding of what goes wrong in the false knowledge ascriptions a little more complex due to the social interpersonal dimension added by telling (which is not thought to be a part of assertion), and there may be relatively fewer philosophical resources about the interpersonal mechanism involved in telling. However, this additional norm crucially explains why contextualists claim that there is a strong intuition both the knowledge ascription by Thelma and the denial of knowledge by Louise are true, or the shiftiness intuition more generally: the intuition that in high-stakes scenarios, more epistemic warrant is needed for permissible assertion, and that in low-stakes scenarios, less

epistemic warrant is needed for permissible assertion. In what follows, I explain the difference between the speech acts of telling and asserting, then argue exactly through what mechanism the interpersonal norm of telling can raise or lower the amount of epistemic warrant required for permissible telling, or even override the epistemic norm of telling.

Section 3.2 – Telling and its norms

This section will be about the speech act of telling, how to understand it, its norms, how it differs from assertion, and finally, why it is more suitable a speech act to understand the cases put forward by the contextualist. Literature on telling has really picked up for the past three decades, thankfully, so there are relatively more resources to draw from than there would have been 20 years ago.

Perhaps due to the very large (almost complete) degree of overlap between the two, it might seem as though nothing of value can be added by refocusing the shiftiness dilemma on telling. However, differences exist between the two speech acts. Telling itself requires two arguments – the person who tells and the person told. Assertion does not. One is able *assert* that p isolated in a room, whereas it is impossible to simply "tell" that p while isolated in a room. That one can assert and fail to tell strongly indicates that the two are not one and the same speech act. This then invites the question of what the norms governing telling are.

A starting point for distinguishing between assertion and telling that is proposed by Charlie Pelling (2014) is to ask what the communicative intentions characteristic of telling are, then ask the same about assertion, and see whether they differ in any way. He claims that the intention of addressing one's audience is present in telling but not necessarily in asserting. Pelling illustrates this by using an example created by David Owens (2006) of a secret diary filled with intimate thoughts. In a case like this, there is an intention not to communicate or

address the content to anyone, but we still want to say that the author made assertions in their diary. Pelling goes on from here to argue that not only does this indicate that telling and asserting are governed by different norms, but uses further thought experiments to argue that it can be permissible for one to make an assertion where it might be impermissible in the same instance to tell someone that very same thing. I think his argument ultimately does not succeed to establish that the knowledge norm does govern testimony as an epistemic norm, as he aims to do, but I take him to be successful in arguing that there is something over and above assertion that is part of the normativity of telling, and that should explain how telling is audience-directed. All that is required for my argument to follow is that telling and asserting are not the same thing, not that the knowledge norm of assertion might not apply to telling. I do not have a preference either way, and it only adds to the resilience of my argument if it works both ways. In what follows, then, I will present the proposals that have been made concerning the normativity present in telling that is absent from assertion. It is clear that these are different, now it is important to see how. I will then show how the additional interpersonal norm can have an effect on the epistemic warrant required in a particular case of telling.

One proposal for what this additional norm of telling is comes from Richard Moran (2005), who claims that telling involves a type of promise from the speaker, an assurance for the truth of what they are saying. He argues that cases of overhearing or observation of behaviour cannot endow the hearer with the assurance that is involved in telling.

"On the assurance view, dependence on someone's freely assuming responsibility for the truth of P, presenting herself as a kind of guarantor, provides me with a characteristic reason to believe, different in kind from anything provided by impersonal evidence alone." (Moran 8)

To add to this, Moran considers a view on which this interpersonal aspect is absent from telling (the evidential view of testimony – treating one's testimony that p as nothing more than evidence of p) and offers a list of reasons why this does not characterise the practice of telling well. One of these is that the speaker's and the hearer's perspective of their common activity clashes by adopting this attitude, with interlocutors being out of sync if the hearer treating one's testimony as a reason for belief only in virtue as its value as evidence. The same can be said if the speaker offers their testimony only as a piece of evidence, thereby trying to remove the assurance of their telling – it makes for an uncommon and uncharacteristic case of testimony ("here is my assertion, do with it what you will"). Moran's point deals with reasons for belief more than the normativity of telling, but for our purposes it is good enough to treat this as a case where the interpersonal norm is absent. It is therefore impossible to construe an account of telling that perfectly mirrors assertion.

In a similar vein, Hinchman thinks that testimony involves an invitation to trust from the speaker to the hearer. More specifically, he argues that in telling, speakers "can be responsible for your belief because in telling you I invite you to trust me, and in inviting you to trust me I take a degree of responsibility for your trust by certifying its status as reasonable." (Hinchman 587).

What we are able to see in both cases is that there is an interpersonal norm governing telling that is absent from assertion, that has to do with the speaker's relation to the information offered to the hearer. This is their utterance offered as a promise, or as a invitation to trust, or as an assurance to the hearer for the truth of p. In any case, we can think of the interpersonal norm of telling as the speaker assuming responsibility for the truth of p to their audience. Since the views are towing a similar line, there is no point in deciding between Hinchman's and Moran's views, or any other views on the interpersonal normativity of telling. Both of these so far share a high degree of similarity, and both are able to supply

the interpersonal condition that is absent from the knowledge norm of assertion. Moreover, my argument will work on both accounts.

Section 3.3 – A more suitable framework for Thelma and Louise

In this section I explain how telling is a more suitable speech act to describe what is going on in DeRose's bank cases, Thelma and Louise cases, as well as Simion's aspirin cases that we have seen in chapter 2.

By proposing that we focus on telling as opposed to assertion, I do not wish to modify any of the examples presented. My proposal is one of providing a more accurate description of what occurs in these cases. Indeed, if I were to modify these examples, I could be accused of changing the contextualist's subject matter, or of not being to apply my analysis to faithful representations of the contextualist thesis. Here are the reasons why my proposal does not amount to modifying or making up different examples than proposed by DeRose or Simion: Telling is a sub-species of assertion, rather than a speech act that is different from it. Secondly, while in theory, assertion applies to more cases (such as asserting things in a journal or asserting things out loud with no audience), in practice, most cases of assertions are cases of telling, while all cases of telling are cases of assertion. However, what is crucial for my argument is that there are no cases that concern the shiftiness dilemma that could apply to assertion but not to telling: the only cases where the two do not overlap are border cases in which there is no audience[6].

[6] Even if the contextualist will wish to include cases in which one speaks to oneself as some kind of proof that assertion is a better framework for the contextualist's example, that will not do. One can be an audience to oneself, which is why the interpersonal norm can apply to cases of talking to oneself. Things are only different when one asserts thing in their sleep, in their journal with no intention of reading it back or thoughtlessly with no intention of 'participating' or receiving what one tells oneself.

In the Thelma and Louise cases, both Thelma and Louise tell their respective audiences that Lena does (or does not) know that John was in the office. In the DeRose's original bank cases (even though I have not described them in this thesis), the speaker *tells* their interlocutor either that they know or that they do not know if a particular bank is open. In the aspirin cases presented in chapter 2, the speaker tells their interlocutor that they have aspirin at home, or to drop by the pharmacy just in case. Apart from clearly being cases in which speakers tell something to their audience, these are cases where stakes are involved, in which it is crucial that the speakers are taking responsibility for the truth of what they are saying. Given that *all* of the cases that will be relevant to the contextualist and invariantist debate will be cases in which there will be practical stakes involved, it is safe to assume that these will be cases in which speakers will tell something to their audience.

In any event, the point of this section is that there is overwhelming reason to believe that telling is a slightly more precise description of the contextualist's imaginary scenarios, and amounts to nothing like a re-imagining but a better look at the same scenarios. If there will be any argument that assertion is indeed more suitable or faithful, the burden of proof will be on the side of the contextualist to explain why the additional cases that apply to assertion are in any way relevant to the shiftiness dilemma.

Section 3.4 – A better warranted assertability manoeuvre

In this section I make my main argument against the epistemic contextualist. To begin with, as we have been able to see, telling is a more suitable framework through which to understand the contextualist's examples. Getting more precise with DeRose's examples allows us to see the picture in more detail, and see what kind of work the interpersonal norm of telling is doing. I do not claim to re-formulate the examples stated by DeRose, or to

change anything about them. In fact, we have seen that all cases of tellings are assertions, so they still very much apply to DeRose's definitions. As well as this, anyone would have a very hard time denying that these are cases of telling.

My argument is that the force behind the shiftiness intuition, and why it strongly seems to anyone who is presented DeRose's examples that in higher stakes scenarios, more epistemic warrant is needed for knowledge ascriptions and denials (or any assertions at all) and conversely for low-stakes scenarios, owes to the interpersonal norm of telling. This norm specifies how much epistemic warrant for p a speaker needs in order to be warranted in taking responsibility for the truth of p, which depends on the stakes involved at the context. We can isolate two types of cases that my argument will be able to explain, that contextualism will not:

A) Cases in which we can say with certitude that the speaker does not know whether p is true, and yet they are completely warranted in telling their audience that p, due to the very low stakes involved and **B)** Cases in which we can say with certitude that the speaker does know that p is true, yet they are not warranted in telling their audience that p, due to the very high stakes.

These situations are not permitted by the contextualist, and can only be explained by the interpersonal norm of telling. The reason **A** cannot be accounted for by contextualism is that knowing that p is a requirement for the permissibility of one's assertion, and that is precisely the aspect missing from this example. The reason **B** cannot be accounted for by contextualism is that the speaker has already met the epistemic standard for knowing that p, and therefore knows that p, yet is still unwarranted in asserting that p. This could not happen on a contextualist account, since the only factor that determines the permissibility of an assertion is whether or not the knowledge norm is met.

Let us see how cases like **A** can happen: imagine here cases in which the speaker has very good epistemic warrant for p, believes that p, but still their justification for believing that p is not enough to qualify them as knowing that p. Simion's Aspirin High case works wonderfully for us in this case: remember that in that example, the stakes are incredibly high. Whether or not the speaker has aspirin at home determines whether or not their sister's baby gets their urgently needed medication. In the Aspirin Low case, a lot less is hinges on the context, as only a minor headache is at stake. Given that I am an invariantist, and that Aspirin High and Aspirin Low involve the same amount of epistemic warrant on the part of the speaker, I only have two decisions ahead of me: either I say that the speaker is Aspirin High is wrong, and they do know very well that there is aspirin at home, or the speaker in Aspirin Low is wrong, and they are driving home without fully knowing that there is aspirin at home. I do not favour one of these over the other, and as I have mentioned before, I am not in a position to judge which of the two is true. Crucially, it does not matter which of the two is true, as I am about to argue that the interpersonal norm of telling is able to explain what makes both of the speaker's tellings warranted, whichever possibility obtains.

Let us then account for both possibilities and consider the first scenario about the epistemic standing of the speakers, in which their epistemic warrant is in both cases enough to meet the invariant standard for knowing that p. This means that, in Aspirin High, the speaker is wrong to think that they do not know that there is aspirin at home. In Aspirin Low, the case is straightforward, since the speaker knows that they have aspirin at home. However, both speakers are warranted in what they tell their hearers. Here is why:

What this example is assuming is that the shiftiness intuition applies, and it is right to do so. It seems intuitively much more permissible to assert "we have aspirin at home" in a low-stakes scenario, than it would in a high-stakes scenario. For the speaker in the high-stakes scenario to be able to tell their audience "there is aspirin at home", they have to be able

to assume responsibility for the truth that there was aspirin at home, *considering* the weight of the practical stakes involved. Even though the speaker does have enough epistemic warrant that they know that p is true, they are not warranted in taking responsibility for the truth of what they know, since the stakes are so high. The reason that knowledge does not suffice for permissibility is just that the stakes are too high: the gravity of the situation in the (highly unlikely) situation that the speaker was wrong about their knowledge simply overrides their knowledge that they have aspirin at home. This is the door through which practical stakes can influence the permissibility of an assertion. The amount of epistemic warrant required for permissible assertions that p is independent from the amount of epistemic warrant required for the truth of p, since the former involves practical stakes and the latter does not. So we can see how the context-sensitivity that comes with practical stakes influences the permissibility of one's telling that p.

Now we will consider the situation in which the speaker in Aspirin High is right in their assumption that their memory of aspirin does not amount to enough epistemic warrant to meet the requirement for knowledge, and the speaker in Aspirin Low is wrong in assuming that they know that they have aspirin at home, given that both speakers have the same amount of epistemic warrant. This coincides with the above-mentioned **B** cases that contextualists cannot account for, in which one does not have knowledge that p, and is still permissibly asserting that p. In this case, the speaker in Aspirin High is doing the right thing in double-checking and going to the pharmacy, since they do not know if they have aspirin at home. On my view, the speaker in Aspirin Low, even if they do not know that p, is warranted in telling their audience that p, given that the stakes are so low. This is another case in which the interpersonal norm overrides the epistemic norm of telling. First of all, even if the speaker does not reach the epistemic standard required for knowledge, they do meet the epistemic standard required to be justified in believing that p. Being justified in believing that p is true

is one example in which one can permissibly tell another that p, even if they fail to know that p. Secondly, given that the gravity of the situation linked to the possibility that the speaker would be wrong is so insignificant, the speaker is warranted in taking responsibility for the truth of p with an amount of epistemic warrant that falls short of that required for knowledge. In this case, one can fail to meet the epistemic standards that p yet still completely permissibly assert that p, since so the stakes are small enough to warrant taking responsibility for the truth of p on the basis of very little. It is important to underline this role of the counterfactual possibility that the speaker could wrong in their belief that p (from their perspective), and the practical stakes tied to those possibility. On my view, the higher the practical stakes involved in the possibility that the speaker is wrong in their belief that p, the more amount of epistemic warrant is needed in taking responsibility for the truth of p. Remember Asprin High: if the speaker was wrong in their belief that they did not have aspirin at home, the consequences would be too dire, so the amount of epistemic warrant required for knowledge that p is not the same as the amount of epistemic warrant required for a permissible assertion that p, in that case. Conversely, the smaller the practical stakes involved in the possibility that the speaker is wrong in their belief that p, the less amount of epistemic warrant is needed for the speaker to take responsibility for the truth of p.

An interesting border appears: when there would be absolutely no practical stakes involved in the situation that the speaker is wrong in their belief, then it is an open question how much epistemic warrant for p they need to be warranted in taking responsibility for the truth of p. Of course, there are moral and other politeness norms to take account of, in the situation that speakers make assertions blindly, or on the basis of very flimsy evidence, but it prima facie strikes me as true that for those things whose truth or falseness have no bearing whatsoever on us, it is fair game what we say. But this need not be part of my larger argument.

There are many ways in which one can permissibly tell their audience that p without knowing that p. One of these is just if the speaker is justified in their belief that p is true, but they do not have enough justification to actually know that p. There are also cases in which very little is at stake in the situation, and the stakes involved in the likelihood that the speaker is wrong in their belief, are likewise relatively insignificant. Here is one example, which I will also use in chapter 3: I can permissibly tell my 5-year nephew that England has 48 counties on the basis on flimsy epistemic warrant, due to the low stakes of the situation, and the low practical stakes in the situation that I would be wrong. The same statement could not be permissibly addressed to the UK's National Statistician at a national census meeting, due to the high-stakes involved in that situation, and the high-stakes in the situation that I would be wrong.

It is important here to note that this position relies on the truth of epistemic fallibilism, which is the thesis that one is able to know that p even when there exists some epistemic chance for them that non-P may be true. This thesis of one's perceived fallibility in knowing that p is not a rogue idea. This is just the anti-skeptical thesis that one is able to know that p even if one does not have the maximal amount of justification they can obtain (Kim 2016). Fantl and McGrath (2009) define it in a very similar way, that is closer to what possibility I am allowing for, which is that one is able to know that p even when there exists some epistemic chance for them that non-P may be true. This seems right to me.

This perceived chance that one is wrong in their belief diminishes, the stronger amount of epistemic warrant one has, but it is likewise important to bear in mind that there are still stronger and weaker degrees of epistemic warrant that one can have even after the invariant standard for knowledge is met. To offer a brief example, consider how knowledge that p on the basis of second-hand testimony that p is possible, but will place the hearer in a

weaker epistemic standing towards p, even if they know that p, to a different hearer, who has direct perceptual evidence that p, and therefore is in better epistemic standing towards p.

It is important to bear in mind that, on my account, there are multiple factors that influence what makes one's telling that p permissible: the weight of the practical stakes at hand, one's epistemic warrant, and the weight of the stakes if one was wrong in their belief that p, which determine what one is or is not warranted in assuming responsibility for the truth of. I argue that what one is able in assuming responsibility for can override the epistemic norm involved in telling.

It is also important to note that my argument does not amount to a pragmatic encroachment view. I do not wish to argue that pragmatic considerations (the practical stakes involved in a scenario) have any influence in what it takes for a subject to actually *know* that p is true. Quite the opposite, I argue that such practical considerations *cannot* have such an effect on what it takes for one to know that p. What I claim is that pragmatic considerations can and do affect what one is permitted to assert or tell their hearer, as they can only take responsibility for the truth of p, as far as the practical stakes allow them to. When the stakes are very high, we are more reticent to take responsibility, and understandably so.

Chapter 4

-Rebuttals of counterarguments-

Section 4.1 – 4 strong objections and their answers

In this chapter, I defend my thesis from the most direct and potentially destructive counterarguments against my case that I am able to construct on behalf of my opponent. Most of these have to do with the application of the speech act of telling to the contextualist's cases. I do not expect to address *all* of the attacks that can be raised against my arguments, but I hope that the most salient and powerful will be touched on by the end of this chapter. The criticisms should be assumed as coming from an adamant contextualist's perspective, attacking my theory:

1) Let's assume you are right, and that in cases in which the practical stakes are higher, the unassertability of knowledge attributions is owed to the fact that speakers in those situations want to avoid breaches of the interpersonal norm involved in telling (e.g. wanting to avoid taking responsibility for the truth of p, or wanting to avoid assuring the hearer that p is the case, due to the weight of the stakes involved) even when the subjects do meet the invariant epistemic standard for knowledge. We can also assume that speakers do additionally have this desire to adhere to the interpersonal norm of telling, and this does the work to explain the unassertability of the contextualist border cases, rather than the subjects not being able to meet a raised epistemic standard associated with higher stakes of knowledge attributions. Even the invariantist is granted all of this, it still does not explain why, as the practical stakes involved in a context are raised, speakers who find themselves in better epistemic positions, *will* be

able to attribute knowledge to subjects in high stakes scenarios, without being in breach of any norms *and* truthfully. This can only be precisely *because* of a speaker's better epistemic position with regards to p that in higher stakes cases, that they manage to attribute knowledge truthfully and without being in breach of norms of telling. Unsurprisingly, this just shows us that when the practical stakes in a scenario are raised, the only thing that is required for a knowledge attribution to be permissible is one's good enough epistemic position that p. On the other hand, the fact that the need for better epistemic positions seems to be required by scenarios with more demanding practical stakes, does not seem to be explained by anything about telling, let alone the interpersonal normativity which supposedly constitutes it. This directly suggests that it is one's epistemic position in relation to p which will determine the truth-conditions of knowledge attributions (or at least assertability, according to some classical invariantists), and not anything to do with the norms involved in telling. To illustrate my point by means of an example, consider this re-adaptation of DeRose's *Thelma and Louise* case:

Unbeknownst to Thelma, after parting ways and before making their way back home, Louise and Lena went in the office to grab an umbrella that they left on a table close to John's office. As they went inside for their umbrella and looked all around the office, they both saw that John was indeed there, and greeted him to make sure that it was him she saw. Louise and Lena both know very well that John was in the office that day, as they both saw him there. Later on, when Louise is stopped by the police, who are trying to establish whether John was in the office as part of a very high stakes case, Louise will tell the police: "Yes, I know that John was in the office that day, I saw him there." When the police ask Louise if anyone else knows if John

was in the office that day, Louise will tell the police "Lena knows that John was in the office today."

This is different to DeRose's example, as I here aim to show that one can, even in very high stakes scenarios, still be successful in truthfully attributing knowledge to a subject, provided that their corresponding epistemic position is good enough. This example shows that, when the practical stakes are very high, speakers can still truthfully tell their audience that "S knows that P", but only when they are in the corresponding epistemic standing, and only *because* of their epistemic standing. This does not happen because of anything to do with the interpersonal norm of telling, but could only be because the reference of "know(s)" in a higher stakes scenario is such that a higher epistemic standard needs to be met. This is how one can explain why, in the modified example, when Louise does have a lot of direct perceptual evidence, she meets the epistemic standard for "know(s)" that is required in the conversation with the police officer. Conversely, when Louise is in a relatively weaker epistemic standing, as she is in the *original* Thelma and Louise case, she does not meet the required epistemic standard for "know(s)" in the conversation with the police officer, so she correctly denies that Lena knew that John was in the office.

The analysis that only makes use of the interpersonal normativity of telling to distinguish between cases in which knowledge ascriptions are warranted and cases in which they are not, on the other hand, cannot account for why the above happens. In fact, this invariant analysis only muddies the cases more. If such an invariantist argues that in the original Thelma and Louise case, Louise wants to avoid breaching the interpersonal norm of taking responsibility for p (where p is "Lena knows that John was in the office today"), or wants to avoid breaching the interpersonal norm of

> assuring the police that p, given how much is practically riding on the truth of p, they cannot also explain why, when Louise's epistemic standing is raised, the knowledge attribution somehow becomes assertable and true. The contextualist story is so much more straightforward and clearer: this can only be *because* Louise's epistemic position towards p has appropriately changed to meet the context's epistemic requirements for the truth of the knowledge attributions. Meanwhile, the invariantist is stumped: they will have to say that in both the original and re-imagined cases, Louise and Lena meet or do not meet the invariant epistemic standard for knowledge attributions. How will the invariantist explain the difference in the assertability of the more informed Louise's assertion (re-imagined case) and the less informed Louise's assertion (original case), if the only thing that is different between the original and re-imagined scenario is precisely Louise's and Lena's epistemic standing?

My response to 1): This is a well-constructed criticism, but unfortunately gets most of the aspects of my proposal wrong. To begin with, the modified example of Thelma and Louise provided by the contextualist, as well as generalized cases in which practical stakes are high yet speakers are able to attribute knowledge successfully and felicitously, are not in any way incompatible with my version of invariantism, combined with my analysis of telling. To remind you, invariantism claims that there is an invariant epistemic standard that must be met for knowledge ascriptions to be true. I make no claim about what this standard is, and how to locate it. So, one simple rebuttal is to bring up the possibility that in DeRose's (DeRose 2009) original *Thelma and Louise* case, where DeRose's main characters do not actually see John, but infer that he is at the office, the epistemic standards required for the attributing knowledge to Lena are not met, and the epistemic standards for 'knowledge' of John being in the office, are not met by any of the three protagonists. Similarly, it is even

more likely that in the modified Thelma and Louise case, where Louise and Lena actually get to see that John is there beyond any reasonable doubt, a much higher standard is met, and Lena and Louise's evidence is simply above the invariant standard of knowledge, while in the original case, it is below. The invariantist is not tied to any thesis that commits them to saying that in both the original and re-imagined case, they have to say the same thing about whether Lena and Louise actually know if John is in the office. That is a misconstrual of invariantism. It might be that the epistemic standard of knowledge attributions is somewhere in between the original and re-imagined case, it might not. Not much hinges on this in any case.

However, invoking this possibility as an answer to the worry presented in 1) cannot do all the work for invalidating it, since my view is not committed to any particular invariant epistemic standard that governs knowledge attributions, only that there is one such standard. I stand by the claim that I am unable to say where this invariant standard is, I am not ready or able to judge examples on a case by case basis, and doing so is unfeasible. Furthermore, doing so would overlook my larger argument, and is not important for our purposes. My rebuttal will have a lot more strength if it will apply to any invariantist interpretation of both my and DeRose's Thelma and Louise cases, including the scenario in which no one (of Thelma, Lena and Louise) meets the epistemic standard for knowing that John is in the office, the scenario in which everyone does, or in which (in the re-invented case) only Lena and Louise do. So, I must be prepared to answer the criticism, even when the original Thelma and Louise case is one in which, the evidence that our protagonists hold does not meet the invariant standard required for knowing that p or attributing knowledge that p. That is when my opponent's argument is the strongest. I can also defend my argument for the original case of Thelma and Louise, interpreted by the invariantist in a way in which the evidence that the three protagonists hold does meet the epistemic standard required for knowledge I defend my argument below.

In the strongest interpretation that can be derived from DeRose's *original* case, where the evidence that Thelma, Louise and Lena all have falls short of the invariant epistemic standard required for knowledge, as an invariantist, I will begin by accepting, of course, that *both* the attribution and the denial of knowledge of the same subject of the same proposition cannot be true, and that Thelma is literally wrong when she says that "She knows, too" about Lena, given that we are leaving from the assumption that her evidence does not meet the 'knowledge' threshold. According to my analysis, Thelma is however, *warranted* in telling her friends in the tavern that Lena knows, because, she is taking responsibility for the truth of P, or assuring her audience of the truth of P, without actually fully knowing that P. What makes her telling warranted is the fact that so little practically hinges on this – there are no repercussions or consequences for Thelma's assurance that P, so she is warranted in her telling her friends at the tavern (falsely), that Lena knows that John was at the office. More importantly though, Thelma is justified in believing that Lena knew that John had been at the office.

On the other hand, there is much more at stake in Louise's case, as the police is investigating a high-stakes murder case. Because of this, Louise telling the police "Lena knows that John is at the office" would not only be false, but also unwarranted, since Louise should not be taking responsibility for the truth of p, or assuring the police that p is true, given how much is riding on p. This would constitute a breach of the interpersonal norm of telling.

I have shown that my analysis can handle individual cases, but it would be better to provide a comprehensive explanation of how being warranted in telling one's audience that p in high-stakes scenarios relates to one's epistemic standing towards p, so that I can alleviate the contextualist's worry expressed in 1). Even though the contextualist's argument is well constructed, they are wrong in thinking that they will be able to show anything by providing

examples of better epistemic positions and showing that these are examples in which knowledge ascriptions are warranted and true. This is all perfectly compatible with invariantism, as having more epistemic warrant will generally be linked to meeting the invariant epistemic standard for knowledge.

My more general reply to the counter-argument, and the better-informed answer, is that the interpersonal norm that is a part of telling cannot be separated from the amount of epistemic warrant or evidence that p that the speaker holds. In fact, we should naturally expect that, the more evidence or warrant that S will have with regards to p, the easier they will find it to take responsibility with respect to p, or to want to assure their audience that p is true. To see why, first consider that even if one already meets the epistemic standard for knowledge, one can still be in a better or worse epistemic standing towards p. For example, consider the difference in epistemic standing between two people who know that p is true. A has very good, direct perceptual evidence for the truth of p, whereas B has third-hand testimonial evidence for the truth of p. Again, both A and B know that p is true, but A is in better epistemic standing towards p. Holding everything equal between A and B, when participating in conversations with people unaware of the truth of p, we will expect that A will be more ready to take responsibility for the truth of p, and to assure others that p is true, having direct perceptual evidence that p is true (and therefore being in a better epistemic standing towards p, compared to B), while B will not be ***as*** ready to accept responsibility for the truth of p, or to assure others that p, given that they only have third-hand testimonial evidence that p is true, and that they are in a worse epistemic standing to A, even if they both know that p is true. This will very largely depend on the practical stakes involved in the scenario, as even with B's weaker body of evidence, there will be plenty of scenarios in which their third-hand testimonial evidence will suffice for one's being warranted in making

a promise that p is true, or being responsible for the truth of p (e.g. being out with one's friends at a tavern).

We can be even more specific about the way in which speakers can present differences in what they are warranted to tell others, depending on their epistemic standing towards p. We will then tie this discussion of different levels of warrant back to their application to knowledge attributions.

As just mentioned, when speakers are very well positioned epistemically towards p, they will be more willing to take responsibility for p, as the stakes involved increase in weight. For instance, as practical stakes increase, A will be willing to take responsibility for the truth of p or promise that p is true, because they are in very good epistemic standing towards p, i.e. they have direct perceptual evidence that p is true. Even as the stakes evolve to the type of police investigation imagined in the Thelma and Louise scenario, someone like A will be able to take responsibility for the truth of p, and assure others or promise them that p is true, due to their very good epistemic standing. However, someone like B, who also knows that p is true, but through weaker evidence, will be less warranted in and willing to take responsibility for the truth of p as the stakes increase. They will be able to adhere to the interpersonal norm involved in telling when not much practically hinges on p, e.g. in a tavern on a night out, and therefore be warranted in telling their audience that p is true. However, as the stakes increase e.g. in a police investigation, the relative weakness of B's evidence will become salient, even if B knows that p is true, and they will no longer be able to make promises to the truth of p. The stakes are just too high – if B makes a promise that p is true or takes responsibility for the truth of p to the police, and p turns out to be false (which is possible, for all B knows), the outcome would be far too serious. On the other hand, B is able to make promises for the truth of p in a tavern, to disinterested friends, and p can turn out to be false, and nobody will be any the wiser. Therefore, we can see that the interpersonal norm

and one's epistemic standing do not exist in a vacuum. When the latter moves, the former does too. This is just an example of a non-epistemic norm that can cause shifts in the amount of epistemic warrant required for proper assertions.

Let us return to the case of Thelma and Louise. The interpretation I am proposing of this case is an invariantist one, in which neither Thelma, Louise, nor Lena have enough evidence to establish that John was at the office, so it is strictly speaking false to say "Lena knows that John was at the office". Even though she happens to be wrong, Thelma is warranted in telling her friends "Lena knows that p", because only a trivial office bet is at stake, and not much hinges on the truth of p, so she is warranted in taking responsibility for the truth of p with her epistemic standing. This can be true even if Thelma does not know that p. Holding the same body of evidence, Louise would be unwarranted in telling the police officers that Lena knows that p, as so much more hinges on it, and she would not be warranted in taking responsibility for the truth of p in good conscience, which is why Louise simply presents her evidence, and tells the police that Lena does not know if John was in.

It does not matter much whether the three protagonists actually do or do not know if John was in, based on the evidence presented, or what *we* (as readers who know that John was at the office) want to say about the strength of their evidence, and whether it actually meets the epistemic standard for knowledge. We may as well be considering the case in which all three have good enough evidence to establish that John *was* at the office, or the re-imagined case presented in criticism 1), where we could even say that Lena and Louise knew that John was in and Thelma did not. In all scenarios, the explanation of why it may strike us as somehow "correct" or "warranted" that Thelma attributes knowledge to Lena, while Louise denies it of her (or ascribes it when she has more evidence), owes to the weight of the stakes involved, combined with what the speakers are warranted in taking responsibility for. As I have argued in chapter 3, we can generalize this to the shiftiness intuition. The appeal of

DeRose's cases, and the intuition that both the denial and the attribution of knowledge of the same proposition can ring true, owes not to any changes of meaning or reference of the word "know" caused by differences in context, but to the way in which the practical stakes influence what the speakers are warranted in telling their audience.

The point of criticism 1) was to attempt to show that adopting the point of view I am adopting i.e. that telling is a more suitable speech act to understand the contextualist's example, will leave the invariantist (me) unable to account for why speakers who are in a better epistemic standing, will be able in a much better position to assert that "S knows that p" seemingly truthfully, whereas when the same speakers are in a slightly lower epistemic standing, they will not be positioned the same way to assert that "S knows that p", and that the only way this could be explained is through the contextualist thesis. My response was to say that one's epistemic position is importantly linked with the interpersonal norm of telling. To be more precise, the better placed with respect to p one is epistemically, the more warranted they will be in assuming responsibility for the truth of p, to promise that p is true, or to assure their audience that p is true.

The second line of defence to this criticism was to say that, indeed, in the re-imagined Thelma and Louise scenario, two of the protagonists do have more evidence than in the original case. A simple explanation for what happens between Louise's denial of Lena's knowledge in the original case, and her affirmation in the re-imagined case could simply be a matter of Lena and Louise meeting the epistemic standard for knowledge in the re-imagined case, while in DeRose's original case, they are below it. Any way we choose to look at it, someone ends up speaking falsely. As long as the invariantist is able to make claims about how all of the assertions in the original and re-imagined cases are warranted knowledge ascriptions or denials, they need not vindicate Thelma and Louise's knowledge ascriptions and denials as literally true. Criticism 1) misses its mark by a long shot.

2) Your argument is no better than any other warranted assertability manoeuvre. You claim that speakers can be "warranted" in telling others "S knows that P", even if they fail to meet the epistemic standards for "knowing" that S knows that P, but you just expect us to take your word for it. How could one lack the epistemic warrant to back their knowledge that p is true, and still be warranted in telling their speaker that p? If that were true, you would also have to say that speakers can even be warranted in lying about the truth of p.

My response to 2): Another valid argument from the contextualist camp. It is a good question to ask what being warranted even amounts to, if one could be warranted in telling their audience something without knowing it. However, the connection between the truth of p and warrant of telling someone that p starts to lose its appeal once we consider how many cases there are when one happens without the other. It is clear that the first can exist without the second. One is not always warranted in saying everything that is true, as there are norms of politeness that govern our conversations. Even if it were true, it would be hugely unwarranted for one to say "Your pie is making me sick, Margaret", so we can see how the first can exist without the second. What about the other way around? As I have argued in chapter 3, there are plenty of cases in which we can see that the two can come apart cleanly. First of all, one can be justified in their belief about p, and still be technically wrong about the truth of p, so their telling that p is also thereby justified. One example of this is Thelma's assertion[7] in the invariantist interpretation where her amount of epistemic warrant did not suffice to qualify as knowledge that p, but it sufficed for her to be justified in her belief, and to tell her friends at the Tavern, that p. That is one example, here is another: one's being warranted in telling things when they

[7] As you have been able to tell, throughout this thesis, I have at times used "assertion" when "telling that", would have suited my point a little better, as it would have matched my argument to the tee. However, due to considerations of readability, "assertion" reads a bit better, as it does not need a "that" after it.

have a low amount of warrant, in very low stakes scenarios: one telling children about stories, that one is not completely justified in believing.

This gets us to another part of the criticism in 2), which is that when we separate the two, one could even be warranted in lying to their audience. There are a lot of problems with this thought. To begin with, the classical analysis of lying, advanced by Augustine in *De Mendacio* and accepted by many other theorists, including as well as many contemporary analyses of lying, a requirement for being able to lie that p is to know that p is not true. That condition is already not met, if the assumption we are leaving from is that the subject does not know that p. It is important not to confuse the truth of not-P with not knowing whether p is true. My claim is that only for the former, can there be situations in which one is warranted in telling the other that p. Knowing that p is not true is entirely different to the cases I am referring to (having just under the standard for knowing that p), and the epistemic standard for knowing that not-p is true is just the same as for knowing that p is true. Once a speaker has information that p is not true, they are indeed in a position where they can lie about p (according to many theories of lying), but I do not concern myself with these examples.

3) On what grounds do you claim that the examples of knowledge ascriptions and denials of the same subject and predicate that are presented in the contextualist examples are not both true? I understand why you think that they cannot *both* be true. But how is it possible to, in the same breath, accept that the shiftiness intuition has a strong intuitive pull, and to also deny that we must vindicate the truth of the speaker's knowledge attributions and denials? We are reading the same examples, that a contextualist imagined, in which both speakers' claims are true claims.

> To invoke arguments made by DeRose (2009), if anything, we have very strong reasons to accept that, in the original *Thelma and Louise* case, both Thelma and Louise's assertions are true. Intuitions to the effect that a knowledge attribution or denial is true are a lot more reliable than intuitions to the effect that a knowledge attribution or denial is false. If an example of a knowledge ascription looks to us to be true, we already have good grounds for believing that it is true. On the other hand, we are on much shakier ground if we use intuitions to the effect that a speaker is wrong in attributing or denying knowledge of a subject. This is what you are asking us to do, as an invariantist, when you are asking us to imagine that the re-imagined Thelma and Louise case is one in which Thelma is attributing knowledge falsely to Lena in the tavern. Must we simply ignore these strong intuitions and adopt an objective attitude whereby we are in a better position to judge the meaning and the truth of speaker's claims in ordinary conversation? This is what we are forced to do in adopting invariantism, where we claim that speakers are attributing or denying knowledge claims falsely. On the other hand, we can explain away the seemingly contradictory claims of both Thelma and Louise, while also managing to respect and preserve their own thoughts about the truth and meaning of what they are saying, not having to adopt a viewpoint where we, as theorists, are in a better position than speakers are about their claims. If your theory about knowledge attributions leads you to outright reject what are clearly true examples of knowledge attributions, it is the theory that needs revisiting, not the examples.

My response to 3): I am not at all convinced by this type of argument. This argument has been made in various places by DeRose (2009), ascribing more reliability to intuitions that knowledge attributions and denials are true, and claiming that intuitions to the effect that knowledge attributions and denials are not true, are not reliable at all. Criticism 3) then asks

us why we do not accept that, in the original example both Thelma and Louise's assertions are true, given that we have such good reason to believe that they are, based on our intuitions.

An immediate response it to say that, if we do not accept that "know(s)" can have multiple referents, which are determined by context, it follows that one also cannot accept the truth of two contradicting statements like "When Thelma says at the tavern that Lena knows that John was at work, her claim is true" and "When Louise tells the police that Lena doesn't know that John was at work, her claim is true". If one has to go, between the intuition the two above claims are true, and the rule of excluded middle, it will have to be the former, since that is something we can be wrong about. Invariantists will not be able to vindicate the literal truth of contradicting statements of knowledge attributions, and I argue that this does not amount to "biting a bullet" in any major way. Denying the literal truth of some speakers' knowledge attributions, while granting that they are still warranted, given the stakes involved, is hardly a shortcoming.

The principal matter of disagreement, then, becomes the issue of increased reliability of the intuitions of the truth of assertions, compared to the reliability of intuitions to the effect that an assertion is false. Let us then unpack the argument that does not offer much credence to intuitions that a speaker's assertion is false.

Here is one place in which DeRose (2009) puts his view across: "intuitions we have that claims are true are generally more trustworthy than are intuitions we might have to the effect that claims are false." (DeRose, 52) This premise ends up playing a very big role in the foundations of DeRose's contextualism. The main reasons that DeRose uses to back up his idea have to do with their naturalness and appropriateness. Speaking about our intuitions that a certain claim is false, where DeRose just refers to claims of the type we have encountered in the usual contextualist examples, he claims that there is a specific danger of mistaking the intuition that there is something that is not *right* about that claim, for the intuition that the

claim is wrong. Therefore, such negative intuitions are much more tentative to be used in informing our theories. DeRose also backs up his claim by saying that appropriately asserted claims are also a lot more likely to be true. "So if a serious assertion is appropriate, that strongly supports the intuition that it is true, at least from the speaker's point of view." (DeRose 2009) Let us then focus on DeRose's reasons for believing that intuitions that assertions are true are more believable: DeRose says that, in asserting an untrue knowledge claim, there are two main explanations of the errors committed to lead to the speaker's wrong assertion: either it is asserted inappropriately, or the speaker was mistaken or had false beliefs about the underlying matters of fact. In the examples he constructed, neither was the case, and the claims are actually made in perfect propriety, so it is highly unlikely that the speaker's claims are false.

There are multiple reasons for which we should be skeptical of DeRose's strong advocacy that we should be ascribing a much larger degree of credibility to positive intuitions about the truth of knowledge ascriptions and denials, compared to intuitions that one's ascription or denial. To begin with, if we assume that S and p are identical between the two following assertions, the content of one's intuition that an assertion that "S does not know that p" is true, is identical to the content of an intuition that an assertion that "S knows that p" is false. There are differences, of course, in what this entails about how the speaker is saying something that is true or not, and this is, of course, an important difference. The information that is preserved between the two types of intuitions, however, is the information about the state of affairs with regards to S, and whether they know that p is true. Abstracting from the speaker and what we are able to say about them, we can reformulate any example of a positive intuition about the truth of a knowledge claim offered by the contextualist that can be used to back their case, into an intuition that a negative knowledge claim is false. Bringing this up is important only inasmuch as we are able to reformulate imaginary examples of

knowledge denials that are true into examples of knowledge attributions which are false. What we lose, is of course, the speaker's relation to the truth of what they are saying, along with the implications of their evidence and the warrant that they have. I make this point not because I think DeRose is overlooking something obvious (which I don't), but because we should be aware that any examples that can be constructed to favour the contextualist (in which the speaker is always right in what they are saying), can be reconstructed in a way in which the speaker is wrong, while in both scenarios, S did not know that p. It is also important to bear in mind that I am not here referring to the further entailments of both types of intuitions, such as that speakers making assertions that are true are "right" in a way in which speakers who make assertions that are "wrong", or what that says about how warranted the speakers are in their assertions.

There is, however, also a possibility that the speakers may speak with perfect propriety, and have no false belief about the underlying matters of fact, *and* that their claims are false. As we have seen, this is largely discounted by DeRose, but it ultimately falls down to the precise limit of the epistemic standards for knowledge attributions. If an invariantist places it just above the standard that all the protagonists meet in the Thelma and Louise case, no one in the scenario can be said to truly know that John was in the office, so Thelma will be speaking falsely in saying "S knows that p". The same can easily be said for skeptical invariantist positions, which will claim that virtually all sentences of the form "S knows that p" will come out as untrue. Still, this does not mean that any of the speakers need to be mistaken about the underlying matters of fact, or that they are asserting inappropriately, but simply that they are not able to meet the invariant epistemic standard of knowledge attributions.

In fact, according to my argument, it is precisely *due to* the speakers making appropriate claims, as opposed to merely literally true ones, with no mistake about the

underlying matters of fact, that speakers can end up speaking literally falsely. I must stress that this is not due to any overemphasis of a politeness norm, but due to following the interpersonal norm of telling, which requires speakers to take responsibility for the truth of p. So, if a speaker claims falsely that "S does not know that p", one particularly important possibility is that the interpersonal norm of telling is overriding the knowledge norm, such that what is appropriate to say, in that context, due to the practical stakes involved, takes precedence over "S knows that p", which is true.

A possible response that DeRose can offer is that the proposed reformulations of assertions of negated propositions are completely irrelevant to his case, given how much the scenario changes from one case to the other. In reformulating our intuitions so that the speaker is saying something false unknowingly, the speaker's beliefs about the underlying matters of fact will also be different, so the two cases are not equivalent. Again, as I argue above, this need not be the case – speakers may have no mistaken beliefs about the underlying matters of fact, and be speaking appropriately, while making false claims. But, I argue, we get precisely to crux of DeRose's point here, i.e. the differences between intuitions that speakers are saying something false versus intuitions that speakers are saying something true. The only mistake that has to occur is the self-belief that the speaker knows that what they are saying is true. That must occur in order for the above to follow, otherwise the speaker would be engaging in irony, sarcasm, lying, or pretence.

I am not alone in having serious doubts about DeRose's not very well substantiated claim that intuitions that claims are right are much more trustworthy than their counterparts. In his (2005) *Knowledge and Practical Interests*, Jason Stanley makes virtually the same point as I have above, saying that DeRose is committing himself to an awkward position through his distinction "between intuitions we have that assertions of negated propositions are true (which are, according to [DeRose] trustworthy) and intuitions that assertions of non-

negated propositions are false" (DeRose 2005, 123). However, this is as far as Stanley goes, so DeRose charges him with not explicitly carrying his claim further and explaining what is uncomfortable about his distinction (2009), so I will try to carry Stanley's criticism further, and indicate the problem with DeRose's position. The immediate problem that DeRose's methodology raises is that, as argued above, assertions of negated propositions that are true can be completely and truthfully reformulated into assertions of non-negated proposition that are false. From our perspective, a true assertion of "S does not know that p is true" is the logical equivalent of an assertion of "S knows that p" that we know is false.

We can therefore deduct that DeRose cannot be making any claims about which intuitions (that an assertion is true) are more reliable based on anything to do the form of false assertions, so he must be making an empirical claim about which of the two types of intuitions are more likely to turn out true, which is how his article reads most naturally. In this case, he is making the empirical claim that those who have intuitions that an assertion is false are more likely to be wrong in their intuition, but he is also making the empirical claim that assertions denying or attributing knowledge are very likely to be true if used appropriately. This is what is what is uncomfortable with DeRose's basis of contextualism – at first, he appears to base his dismissal of the use of negative intuitions on a theoretical basis, but it becomes clear that his strong preference over the type of intuition that luckily inform contextualist theories is based on an empirical notion of which type of intuition is more likely to actually come out true. This source of discomfort that Stanley points to could perhaps be somewhat alleviated through an empirical backing, which DeRose does not offer. The upshot is that this theoretical need to vindicate the truth of speaker's claims that is such a large part of the basis of epistemic contextualism, is not there, and should not be kept as the standard against which one builds their theory of knowledge attributions.

There is another criticism that is part of 3) that we must address. This is the criticism that by not "vindicating" all speakers' literal knowledge attributions and denials, we are somehow adopting an objective attitude, where we claim, as theorists, to be in a better position to judge the meanings and truth values of speakers' knowledge attributions. This amounts to having a better idea of what speakers actually "mean" and whether they "right".

This is not a criticism made by any contextualists that I have been able to find so far, but one that I have created on their behalf. Its strength lies in the suggestion that the invariantist is inadvertently but problematically also claiming that speakers don't have a good idea of what they mean, or they are so often not "right" in what they tell others they know. This is a valid worry, since it downplays the speakers' own ideas about the truth of their knowledge attributions and denials.

My first response is to push back on the idea that there is some discounting of the speakers' own intuitions concerning their knowledge claims, which is involved in my argument that, in certain scenarios, some speaker's knowledge claims will come out as literally false. Given that there is a categorical denial of the possibility of two statements like Thelma and Louise's being simultaneously true, it is a matter of entailment that one of these will come out false. Moreover, showing the way in which all speakers in the examples given are warranted in their assertions is precisely the opposite of discounting or undervaluing speaker's intuitions. We are able to explain why speakers are warranted in what they tell their audience, validating their reasons behind telling their audience that "S knows that p". Moreover, my position is not adopting any kind of "objective" attitude where the speakers own intuitions do not inform anything about the truth of their own claims. In the very way that the Thelma and Louise examples are constructed, the reader is put in the position of simply having more information about the scenarios than the protagonists themselves do, and

we are able to see, before the protagonists do, whether their beliefs are true, and whether John is actually in the office.

Consider someone who finds herself in a low-standards context, such as Thelma at the tavern, in a situation in which she has very good reason to believe that p is true, tells her audience "S knows that p", and p is true. The speaker is moreover justified in her belief (as Thelma had very good reason to believe "John was at the office") and in the belief that S knows that p (in Thelma's case, that "Lena knows that John was at the office"), but turned out to be wrong that S knew that p, since S (Lena) did not meet the epistemic standard for knowledge (we are counting this as a possible invariantist interpretation of the original Thelma and Louise case). One way in which a person can be warranted in telling others "S knows that p", without meeting the invariant standard for knowledge, is for them to only be justified in believing that "S knows that p", without having enough epistemic warrant to reach the standard for knowledge.

Consider now a separate situation, one in which our speaker says "S knows that p", the speaker believes that S knows that p, but does not know "S knows that p", she is not very well justified in believing that S knows that p, and p is not true. There is still room for the speaker being warranted in telling her audience that "S knows that p", but crucially only when the practical stakes involved in the scenario are low enough, and the speaker is warranted in taking responsibility for the truth of p. To appreciate how this can happen, consider two friends sitting on a porch, having long carefree discussions about the government, over some beers. As they get to talking about politics, they reach the subject of surveillance, and one tells the other, about the National Security Agency "I am telling you, they even know where you plan on going on holiday, that's how much they know about you."

Let's say that the speaker making this claim has some justification in believing what they are saying, perhaps due to them following sources reporting on the surveillance tactics

employed by the NSA, but these sources are not always factually correct, so he does not meet the invariant epistemic standard for knowledge attributions with his epistemic warrant. Let us also assume that his claim is false, and that the NSA did not know where the speaker's friend was planning on going on holiday, and the NSA do not track or have data on that kind of personal information about its citizens. In this situation, as I have constructed it, the speaker is still warranted in telling what he did to his friend. I argue, that in this case, it is not merely a matter of the speaker having *enough* justification for them to be *warranted* in saying "S knows that p" but not for actually knowing that S knows that p. It is more than the epistemic normativity behind telling that provides the speaker warrant, it is the low practical stakes involved in the context of the utterance that lower the epistemic standard of what the speaker is warranted in taking responsibility for. To see that it is this additional interpersonal component, and not the epistemic normativity, which explains why the speaker is warranted in telling his friend tales of surveillance, consider how it would no longer warranted to say the same thing to an FBI official, asking the speaker if anyone has any information on where his friend intended to go on holiday. In this case, given how much rides on the truth of p, the speaker is no longer warranted in saying what he told to his friend over a beer. The second way, therefore, through which one can be warranted in telling their audience that S knows that p, is through the interpersonal normativity of one's assertion, and how this is affected by the practical stakes involved in her context of utterance.

We are now able to see that my account of what it takes for one being warranted in *telling* their audience that "S knows that p" is two-fold. There is the component that involves the epistemic normativity behind the belief of the speaker, i.e. how much epistemic warrant one has for believing that p, and whether that meets the invariant standard for knowledge attributions. There is also the component of the interpersonal normativity determined by the practical stakes present at the context of utterance, which specifies the amount of epistemic

warrant that one needs for p, in order to be warranted in taking responsibility for the truth of p. For instance, the amount of evidence that I need to be warranted in telling my 5-year old nephew that England has 48 counties is much lower compared to the evidence that is needed to be warranted in telling the same information to the UK's National Statistician at a national census meeting. At the other extreme, if the stakes are too high, one may even be able to meet the invariant standard required for knowledge, but not be warranted in telling their audience that p, given the weight involved. As mentioned in chapter 3, recall that the Aspirin cases introduced by Simion (2021) are good examples of this. Even if I know that I have aspirin at home, the catastrophic consequences of the very unlikely possibility I am wrong, will outweigh the amount of epistemic warrant I have. This shows that the interpersonal normativity component can override the epistemic normativity component when it comes to telling, if the stakes are high or low enough. Similarly, the former can also override the later, in cases where there is enough epistemic warrant. So even if the contextualist will say that one cannot be warranted in telling their audience that p, if they do not know that p, I am able to retort this by saying that the interpersonal norm can and does over-ride the epistemic norm involved in telling, and vice-versa.

My second response is to say that, even if we were in a position where we had to simply deny that speakers are wrong in their knowledge attributions or denials, without making any claim about what they were warranted in saying, this does not amount to some kind of immediate failure of the argument. Admitting that some speakers may be wrong in their knowledge ascriptions or denials is not a deal-breaker of invariantist theories – speakers happen to be wrong about the literal truth of their assertions frequently, especially when the low stakes they find themselves in allow them to do so without repercussion. To give an example that mirrors the way this happens with knowledge ascriptions and denials, consider how speakers are often literally wrong in saying "I know that everything is going to be ok" in

situations in which the speakers does not know if everything will turn out ok. Whether it is wishful thinking or an erroneous judgement of one's own epistemic standing that gives one cause for making such an assertion, such speakers are very frequently warranted in saying what they do due to a multitude of possible reasons, and it would be indeed problematic to adopt a position in which a speaker's own intuitions are less valuable than an epistemologist who thinks they have a good idea of where knowledge starts and where it ends. Even so, this line of argument is one I have no need to fall back on, since I show how speakers are warranted in what they tell their audience. My argument entails that one need not know that p is true, to be necessarily warranted in telling their audience that p, as long as the practical stakes allow them to tell others that p on a weaker basis than knowledge. It also entails that one can know that p is true, and still not be warranted in telling their audience that p, if the practical stakes are high enough, and that the invariant standard for knowledge is not the same as the standard for epistemic warrant for those high stakes.

The upshot of this response, is that by not "vindicating" the literal truth of all speakers' knowledge claims and denials, I am not exposing my position to any major drawback, given that I present a strong argument that all speakers are still warranted in their knowledge attributions and denials, and I explain how this happens. The requirement for vindicating the truth of speakers' knowledge claims is a *very* important part of the "ordinary language basis" (DeRose 2009) of contextualism, on the other hand, and we have seen that there is no such theoretical need when analysing speakers' knowledge claims, and we have also seen that, as we explain how speakers are warranted in telling their audience that "S know(s)/does not know" we are not minimizing the speaker's own intuitions about the truth of what they tell their audience in knowledge attributions or what they mean by them. In dispelling argument 3), we manage to simultaneously protect my version of invariantism from charges of undervaluing the speaker's intuitions about their knowledge attributions, and

we also manage to remove a significant portion of the basis of the contextualist's argument. Not only is counter-argument 3) null, since the assumptions on which it is based are wrong, but it also shows the weak foundations of epistemic contextualism, which is partially kept up on this very controversial claim about the reliability of positive intuitions, compared to negative ones. Below is the last criticism I refute:

4) You cannot just take the imaginary scenarios constructed by someone else, that are presented as cases of assertion, and suddenly claim that they are something else, as cases of telling. These have been constructed carefully, as examples of assertions of knowledge attributions, and modifying the examples amounts to referring to different examples altogether. If you want to talk about the speech act of telling, then go ahead. But contextualism of knowledge attributions and invariantism, respectively, are views that concern assertion, and the shiftiness dilemma is one that is specific to assertion. Even if you do manage to construct your own warranted assertability maneuver, it is one that can only make a point about particular cases of telling, and that does not suit the larger conversation had between the contextualist and the invariantist.

My response to 4): This is another strong counter-argument, but there are even stronger reasons to resist its tempting reasoning. The first reason it does not succeed in creating a chink my view's armour, is that it is mistaken to view assertion and telling as two completely different speech acts. The reason that telling has even been overlooked in this whole debate, is explained by the fact that the two speech acts are so similar. But I do not only base my argument on their similarity, but on their belonging together: telling, as I have mentioned, is a sub-species of assertion. This means that through insisting that we analyze the contextualist's examples as cases of telling, I am not denying that those examples are cases of assertion. In fact, I agree with the contextualist, these *are* cases of assertion, only they are not *just* cases of assertion, they are cases of speakers asserting things to each other,

taking responsibility for the truth of what they are saying, in a conversation. That narrows it down, and there is value to be found in being more specific, and indicating that these are cases of telling.

Another point to add is that I am not "making something up" by getting a little more specific about these examples and saying that they are cases of telling. In fact, it is almost impossible to think of two interlocutors in a conversation, as presented in DeRose's examples, who assert things to each other, yet do not tell each other things. This only indicates that they could not be *only* cases of assertion, and that I am not adding anything to these examples, just making sure we understand the full extent of what is happening in them. This also shows that my proposed honing in on the examples does not amount to a change to what was presented initially. So again, I need not even disagree with the contextualist on this point, I merely wish to be more precise about what is happening in their imaginary scenarios. So, concerning the first part of the criticism, I am not claiming that they are something else than presented, but being more specific about what kind of assertions they are.

The second criticism launched in 4) is that all of the examples presented in the contextualist and invariantist debate, that the criticism dilemma talks about, are examples of assertion, and that even if I do make a point, it is a minor one, that does not apply to the main debate. This is tempting to follow, and I would first say that the scale of my argument is only as big as the amount of cases it applies to. However, I want to say something stronger: that it applies to *all* of the cases that pertain to the shiftiness dilemma, and that all of the examples that would have any relevance to the contextualist and invariantist debate can *only* be cases of telling. This is because 4) would only have any power if there is another subset of assertions that are not cases of telling, to which the shiftiness dilemma applies. This can only involve cases of simply asserting things on one's own, making pretense assertions (in a play), etc. Even if one would try to use this argument that the shiftiness dilemma applies to more cases

than just telling, because it applies to asserting things on one's own, it would not amount to a good argument at all, as it would only be small difference that is not on anyone's mind when thinking about the contextualist and invariantist debates. As it happens, it does not even apply to making assertions blankly, because one can also tell themselves things, and it is a much more appropriate way to understand what happens in talking to oneself as opposed to assertion. It also involves taking responsibility for the truth of what one is telling oneself, which is why it is a better framework for understanding this potential addition.

For the imaginary cases created by DeRose to work to do any intuitive heavy lifting at all (which is important to contextualism), these need to be cases of telling, as these only have intuitive power if they are cases in which speakers take responsibility for the truth of p, promise that p is true, or assure their audience for the truth of p. It is the emphasis of the practical stakes, and its importance in the conversation, that does not allow us to abstract from the interpersonal norm involved in the speaker's telling, and forces us to look at these cases as cases of telling. So, concerning the second part of 4), telling applies to all the cases in the contextualist and invariantist debate, and these can *only* apply to the speech act of telling.

Section 4.2 - Conclusion

In this thesis I have tried to defend an invariantist position about knowledge attributions, arguing that the examples motivating the contextualist case are too vague in their description. I then proposed my own, more accurate understanding of the contextualist's imaginary scenarios. The purpose of this was to argue that the contextualist scenarios support a different analysis to the one proposed by contextualists, rejecting the thesis that the reference of knowledge attribution can vary from one context to the other. By doing this, I

also manage to reject the shiftiness dilemma – i.e. the contextualist argument that one needs to make a choice between endorsing the knowledge norm of assertion and classical invariantism – I have showed one way in which the two *can* coexist.

I made this case partly by maintaining that speakers don't *purely* assert things at each other – they tell each other things. This apparently pedantic point has an important consequence: it means that in engaging in this speech act (telling), speakers are subject to interpersonal norm(s) that are not constitutive of assertion. I then showed some of the ways in which the interpersonal norm that governs telling has been conceived: speakers making promises to hearers about the truth of p, taking responsibility for the truth of p, or assuring the hearer that p is the case. Indeed, contextualists are right to say that our intuitions show that knowledge attributions and denials are used differently depending on the practical stakes involved in a context. However we will choose to interpret the interpersonal norm, paying attention to it when re-evaluating DeRose's examples will highlight the real source of the context-sensitivity of knowledge attributions. The interpersonal norm of telling helps us see that it is not the *reference* of the knowledge claims and denials that shifts along with their appropriate uses. The reference of knowledge claims and denials stays tied to an invariant epistemic standard that governs knowledge. What changes with context-sensitive knowledge attributions and denials is how warranted speakers are in taking responsibility for the truth of what they are telling their audience, which varies along with practical stakes. In high-stakes scenarios, speakers may not be warranted in taking responsibility for the truth of p, even when they know that p is true. In low-stakes scenarios, speakers may be warranted in taking responsibility for the truth of p, even when their amount of epistemic warrant amounts to very little and they do not meet the epistemic standard required for knowledge, as not much rides on their telling. I argued that what plays a significant role in the appropriateness or warrant of one's telling is the gravity and possibility of their being wrong. In high-stakes

scenarios in which one knows that p, even though one's own idea of their likelihood of being wrong is very small (but not non-existent), the gravity of the possibility may be so large that it would make one unwarranted in telling their audience that p is true.

This was my argument against the context-sensitivity of the reference of knowledge attributions, in favour of an invariant epistemic warrant. Equipped with this, I argued that the invariantist need not choose between their position and the knowledge norm. I then also looked at the ordinary language basis of DeRose's contextualism, arguing that it fails to provide a good theoretical basis for contextualism.

In chapter 1, I gave an extensive overview of epistemic contextualism, its history, the relevant-alternative theories it stemmed from, and invariantist theories. I also provided a neutral and terminology-light explanation of epistemic contextualism, based on Kaplan's distinction between content and character. I then introduced the reader to the normativity of speech acts, so I can give an adequate explanation of the knowledge norm of assertion, and its role in the invariantist-contextualist debate.

In chapter 2, I presented the shiftiness dilemma in epistemology, using the information introduced in chapter 1. I showed how the knowledge norm of assertion can be used to argue that classical invariantism is incompatible with KNA. According to the contextualist, we either introduce the variability that explains the context-sensitivity of knowledge attributions through our theory of knowledge, or we have to introduce context-sensitivity within the knowledge norm of assertion. The first option rules out classical invariantism, and the second option rules out adopting the knowledge norm of assertion. I have then presented arguments that have been made against the shiftiness dilemma, showing the strengths and weaknesses of both sides of the debate.

In chapter 3, I made my argument against contextualism, vouching for a more adequate analysis of the contextualist that looks at telling as the speech act common to the

scenarios stipulated by the contextualist, as opposed to assertion. Then, I explained how telling can explain context-sensitivity of knowledge attributions and denials with practical stakes, through its interpersonal norm, i.e. the speaker taking responsibility for the truth of p, in telling his audience that p. I showed how this can override the epistemic norm of telling, and how the epistemic norm can, in turn, override the interpersonal norm. Specifically, it is possible for one's telling that p to be unwarranted even if one knows that p, if the practical stakes are too high, due to the gravity of the stakes overriding one's epistemic warrant. Conversely, when one's epistemic warrant is especially low, low practical stakes and their lack of gravity can offer warrant to one's telling that p, even when they do not meet the invariant epistemic standard for knowing that p.

In chapter 4, I finished by responding to the strongest arguments that can be brought against my thesis. I responded to criticisms that assertability in high-stakes scenarios is restored by raising one's epistemic warrant, which would affirm contextualism and refute invariantism. I responded to criticisms that my argument amounts to just another warranted assertability manoeuvre, and that my grounds for arguing that one's knowledge claim can be warranted but untrue, are shaky. I also answered the attack that our intuitions of the truth of knowledge claims and denial cases are heavily on the side of the contextualist, and that not accounting for them is a mark against my theory. I finally refuted the counter-argument that my insistence to re-analyse the contextualists cases is too vague and tenuous: that it is not up to me to re-analyse the contextualist examples, that the difference between asserting and telling is non-existent, and that if it does exist, my conclusion only applies to telling and is therefore irrelevant to the larger debate. I replied by arguing that the contextualist could not object to the more precise speech act analysis, given that it is suited to all of their scenarios, importantly also applies to all examples that have relevance to the debate, and interest

contextualists and invariantists, and gave further reasons for the differences between telling and asserting.

I end this thesis by driving home the point that we do not assert things in a vacuum to each other, and that we are responsible for what we *tell* each other. At times, our restraint from taking the responsibility of telling our interlocutor something may look like it amounts to falling short of us meeting some epistemic norm. Really, sometimes we are just sensitive to the possible effect of the very high stakes, and want to align ourselves with the seriousness of the situation. Recall a very serious conversation with a person you care about, when the stakes where very high. Were you really prepared, in such an important scenario, to take responsibility for the truth of everything you know that was relevant of the situation? Why not?

Bibliography

Adler, Jonathan E. 2012. “Contextualism and Fallibility: Pragmatic Encroachment, Possibility, and Strength of Epistemic Position.” *Synthese* 188 (2): 247–72.

Bach, Kent, 'Knowledge In and Out of Context', in Joseph Keim Campbell, Michael O'Rourke, and Harry S. Silverstein (eds), Knowledge and Skepticism (Cambridge, MA, 2010)

Benz, Anton. 2021. “Epistemic Perspectives and Communicative Acts.” *Frontiers in Communication* 6. https://doi.org/10.3389/fcomm.2021.612733.

Blaauw, Martijn. 2008. “Subject Sensitive Invariantism: In Memoriam.” *The Philosophical Quarterly* 58 (231): 318–25.

Braun, David. 2017. “Indexicals.” In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Summer 2017. Metaphysics Research Lab, Stanford University.

Brown, J. 2006. “Contextualism and Warranted Assertibility Manoeuvres.” *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 130 (3): 407–35.

Brown, Jessica. 2008a. “Subject-Sensitive Invariantism and the Knowledge Norm for Practical Reasoning.” *Noûs* 42 (2): 167–89.

Cohen, Stewart. 1999. “Contextualism, Skepticism, and the Structure of Reasons.” *Philosophical Perspectives. A Supplement to Nous* 13: 57–89.

DeRose, Keith. 1992. “Contextualism and Knowledge Attributions.” *Philosophy and Phenomenological Research* 52 (4): 913–29.

DeRose, Keith. 1996. “Relevant Alternatives and the Content of Knowledge Attributions.” *Philosophy and Phenomenological Research* 56 (1): 193–97.

DeRose, Keith. 2002. “Assertion, Knowledge, and Context.” *The Philosophical Review* 111 (2): 167–203.

DeRose, Keith. 2009. *The Case for Contextualism: Knowledge, Skepticism, and Context, Vol. 1*. Oxford University Press.

Douven, Igor. 2006. “Assertion, Knowledge, and Rational Credibility.” *The Philosophical Review* 115 (4): 449–85.

Dretske, Fred. 1970. “Epistemic Operators.” *The Journal of Philosophy* 67 (24): 1007–23.

Fantl, Jeremy, and Matthew McGrath. 2002. “Evidence, Pragmatics, and Justification.” *The Philosophical Review* 111 (1): 67–94.

Fantl, Jeremy, and Matthew McGrath 2009. “Knowledge in an Uncertain World,” November.

G. C. Stine. 1976. “Skepticism, Relevant Alternatives, and Deductive Closure.” *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 29 (4): 249–61.

Gerken, Mikkel. 2012. “Discursive Justification and Skepticism.” *Synthese* 189 (2): 373–94.

Goldman, Alvin I. 1976. “Discrimination and Perceptual Knowledge.” *The Journal of Philosophy* 73 (20): 771–91.

Greenough, Patrick. 2011. “Truth-relativism, Norm-relativism, and Assertion.” In *Assertion*, 197–231. Oxford University Press.

Hambourger, Robert. 1987. “Justified Assertion and the Relativity of Knowledge.” *Philosophical Studies* 51 (2): 241–69.

Hazlett, Allan. 2009. “Knowledge and Conversation.” *Philosophy and Phenomenological Research* 78 (3): 591–620.

Hazlett, Allan. 2010. “The Myth of Factive Verbs.” *Philosophy and Phenomenological Research* 80 (3): 497–522.

Hermanowicz, Erika T. “Augustine on Lying.” Speculum, vol. 93, no. 3, 2018, pp. 699–727.

Heller, Mark. 1989. “Relevant Alternatives.” *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 55 (1): 23–40.

Hinchman, Edward S. 2005. “Telling as Inviting to Trust.” *Philosophy and Phenomenological Research* 70 (3): 562–87.

Kaplan, David. 2009. “Between Character and Content and Its Applications.” 2009.

Kaplan, David. 2009 “An Essay on the Semantics, Logic, Metaphysics, and Epistemology of Demonstratives and Other Indexicals.”

Kim, Brian. 2017. “Pragmatic Encroachment in Epistemology.” *Philosophy Compass* 12 (5): e12415.

Lewis, Max. 2021. “The Knowledge Norm of Assertion: Keep It Simple.” *Synthese* 199 (5-6): 12963–84.

Marsili, Neri, and Alex Wiegmann. 2021. “Should I Say That? An Experimental Investigation of the Norm of Assertion.” *Cognition* 212 (July): 104657.

Moran, Richard. 2006. “Getting Told and Being Believed.” *The Epistemology of Testimony*, 272–306.

Myers-Schulz, Blake, and Eric Schwitzgebel. 2013a. “Knowing That P without Believing That P.” *Nous* 47 (2): 371–84.

Myers-Schulz, Blake, and Eric Schwitzgebel. 2013b. “Knowing That P without Believing

That P.” *Nous* 47 (2): 371–84.

Owens, David. 2006. “Testimony and Assertion.” *Philosophical Studies* 130 (1): 105–29.

Pelling, Charlie. 2013. “Assertion and The Provision of Knowledge.” *The Philosophical Quarterly* 63 (251): 293–312.

Pelling, Charlie. 2014. “Assertion, Telling, and Epistemic Norms.” *Australasian Journal of Philosophy* 92 (2): 335–48.

Radford, Colin. 1966. “Knowledge: By Examples.” *Analysis* 27 (1): 1–11.

Rysiew, Patrick. 2021. “Epistemic Contextualism.” In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Spring 2021. Metaphysics Research Lab, Stanford University. https://plato.stanford.edu/archives/spr2021/entries/contextualism-epistemology/.

Sbisà, Marina. 2018. “Varieties of Speech Act Norms.” In *Normativity and Variety of Speech Actions. Edited by M. Witek and I. Witczak-Plisiecka. Poznań Studies in the Philosophy of the Sciences and the Humanities , Vol. 112.*, 23–50. unknown.

Simion, Mona. 2021. “Shifty Speech and Independent Thought: Epistemic Normativity in Context,” April. https://doi.org/10.1093/oso/9780192895288.001.0001.

Speaks, Jeff. “Transparency, Intentionalism, and the Nature of Perceptual Content.” Philosophy and Phenomenological Research, vol. 79, no. 3, 2009, pp. 539–573.

Speaks, Jeff. 2021. “Theories of Meaning.” In *The Stanford Encyclopedia of Philosophy*, edited by Edward N. Zalta, Spring 2021. Metaphysics Research Lab, Stanford University.

Stanley, Jason. Knowledge and Practical Interests. Oxford University Press, 2005.

Stine, G. C. “Skepticism, Relevant Alternatives, and Deductive Closure.” Philosophical Studies, vol. 29, no. 4, 1976, pp. 249–261.

Stone, Jim. 2000. “Skepticism as a Theory of Knowledge.” *Philosophy and Phenomenological Research* 60 (3): 527–45.

Unger, Peter. “A Defense of Skepticism.” The Philosophical Review, vol. 80, no. 2, 1971, pp. 198–219.

Unger, Peter K. Ignorance [Electronic Resource]: a Case for Scepticism. Clarendon Press, 1975.

Unger, Peter K. Philosophical Relativity. University of Minnesota Press, 1984.

Watson, Gary. 2004. “Asserting and Promising.” *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 117 (1/2): 57–77.

Williamson, Timothy, ' Knowledge and Scepticism', in Frank Jackson, and Michael Smith

(eds), The Oxford Handbook of Contemporary Philosophy (2007; online edn, Oxford Academic, 2 Sept. 2009),

Williamson, Timothy. Knowledge and Its Limits. Oxford University Press, 2000.

Williamson, Timothy 2005. “Contextualism, Subject-Sensitive Invariantism, and Knowledge of Knowledge.” *The Philosophical Quarterly* 55 (219): 213–35.