



MASTERARBEIT | MASTER'S THESIS

Titel | Title

Deception, the Lying-Misleading Distinction, and Strategic
Subject-Change

verfasst von | submitted by

Gabriel Levc BA

angestrebter akademischer Grad | in partial fulfilment of the requirements for the degree of
Master of Arts (MA)

Wien | Vienna, 2024

Studienkennzahl lt. Studienblatt | Degree
programme code as it appears on the
student record sheet:

UA 066 941

Studienrichtung lt. Studienblatt | Degree
programme as it appears on the student
record sheet:

Masterstudium Philosophie

Betreut von | Supervisor:

Univ.-Prof. Paulina Sliwa PhD

Abstract

Das Ziel dieser Arbeit ist es, einen Beitrag zu aktuellen Entwicklungen in der philosophischen Auseinandersetzung mit dem Thema *sprachlicher Täuschung* (linguistic deception) zu leisten. Ich werde mich hierbei auf zwei spezifische Themenkomplexe konzentrieren. Der erste betrifft den moralischen Status der Unterscheidung zwischen *Lügen* einerseits und *Irreführung* andererseits (lying-misleading-distinction). Ich argumentiere, dass entgegen der vorherrschenden Meinung Irreführen in manchen Fällen moralisch negativer zu bewerten ist als Lügen, da Irreführen die getäuschten Personen in eine schlechtere epistemische Position bezüglich ihrer eigenen Täuschung bringt. Des Weiteren zeige ich auf, dass diese schlechtere epistemische Position in vielen Fällen moralisch relevant ist.

Das zweite Thema betrifft die Methoden, mit denen Sprecher*innen die Konsequenzen ihrer problematischen Äußerungen vermeiden. Ich untersuche hierzu eine in der philosophischen Literatur bisher wenig beachtete rhetorische Strategie, die des *strategischen Themenwechsels* (strategic subject-change). Ausgehend davon argumentiere ich, dass im Vergleich zur *plausiblen Abstreitbarkeit* (plausible deniability), die derzeit meist als Erklärung dafür dient, dass Sprecher*innen Schuldzuweisungen für normverletzende Äußerungen vermeiden, strategische Themenwechsel besser geeignet sind, um derartige Fälle zu erklären. Anschließend diskutiere ich einige Konsequenzen, die das Phänomen des strategischen Themenwechsels für allgemeinere philosophische Debatten über sprachliche Täuschung hat.

Acknowledgements

I would be lying (or perhaps misleading) if I didn't say that this was the part of writing a thesis that I was looking forward to the most. I am deeply indebted to a lot of amazing people for being able to complete this Master's thesis, and my philosophy degree more generally. There are probably too many to mention all of them, but I will try anyways.

First and foremost, I want to thank my supervisor Paulina Sliwa for supporting this thesis throughout the entire mess it ended up being. I am deeply grateful for all the discussions we've had, the advice she has given me, and her encouragement. Without her, my writing would be even wordier than it is now. I am also grateful to Katharina Felka for inciting my love for the philosophy of language and tolerating the many questions this sparked, to Dirk Kindermann for inciting my love for social/political philosophy of language and tolerating the many questions this sparked, and to Indrek Reiland for his support, especially with the third chapter of this thesis.

My deepest gratitude goes to my mother, Barbara Levc. Through these last couple of years (and, really, for my whole life), she has supported me through every decision, helped me overcome every obstacle, and has celebrated every success with me. This thesis would not have been possible without her, both in a literal and figurative sense.

I thank my grandparents, Peter and Christa Levc, for their encouragement and love, and my father, Erich Levc-Münzer for his unwavering support. I would also like to apologize for the early seminar papers that he so generously proofread.

I also want to express my heartfelt gratitude to Ella Berger and Jonas Piringer, for their endless help, kind words, and love throughout the last years; I do not know where I would be without them. Likewise, I am deeply indebted to Cora Wiedmann for her being the best possible friend, for all her silliness, kindness, and support. Throughout all the weird things that happened in the last couple of years, she has been nothing but a constant source of happiness.

Special thanks go to Matt Dougherty and Martin Niederl who have generously proofread parts of this thesis, sometimes on very short notice, and whose thoughtful comments and criticism have helped me a great deal.

This thesis and everything that preceded it would not have been possible without the encouragement, advice, support and friendship of many people. In addition to the people listed above, I want to thank Paul Astner, Jan Balcerowski, Jakob Bock, Cheyenne Brock, Anna-Maria Edlinger, Bailey Fernandes, Lena Golker, Eva Hijlkema, Alina Jacobs, Sophie Joch, Yara Katnik,

Anna Klieber, Mirela Koleva, Reid Kurashige, Veronika Lassl, Jakob Lissy, Paula Majcenovic, Charlotte Meyer, Martyna Meyer, Ida Miczke, Jeanne Pang, Pauline Paulik, Veronika Rom, Felix Schalk, Michael Stöger, and Marlene Valek for everything, really.

Furthermore, I want to thank the participants of the SoPHIA 2022 and the Workshop “The Foundations of Propagandistic speech” for their helpful feedback on some of the ideas that went into the first part of this thesis.

Finally, I am deeply grateful to the Vienna Forum for Analytic Philosophy and all its members. Thank you for giving me a philosophical home in Vienna.

Table of contents

Abstract.....	1
Acknowledgements	2
1. Introduction	6
2. Literature overview: Current developments in the philosophy of linguistic deception.	8
2.1 Linguistic devices used for deception.....	8
2.2 The moral status of deception	12
2.3 Deception and Deniability	15
3. Lying, Misleading and Epistemic Advantage	19
3.1 Introduction	19
3.2 The lying-misleading distinction	20
3.3 Arguments for lying being worse than misleading	21
3.3.1 The Kantian argument.....	21
3.3.2 Berstler: Lying is unfair in the face of conventions	25
3.4 An argument for misleading being worse than lying.....	26
3.5 The epistemic consequences of misleading	28
3.6 Conclusion.....	31
4. Strategic Subject-change, plausible deniability, and dodgy politicians.....	32
4.1 Introduction	32
4.2 What is strategic subject-change?	33
4.2.1 What is strategic subject-change?	33
4.2.2 Why is strategic subject-change successful?.....	35
4.3 Strategic subject-change and plausible deniability	38
4.3.1 Plausible and not-so-plausible deniability	39
4.3.2 How strategic-subject change can help.....	40
4.4 Strategic subject-change in political speech.....	44
4.4.1 Politicians changing the subject: Some cases	44
4.4.2 Why do politicians strategically change the subject?	45

4.4.3 Why strategic subject-change is successful in political contexts.....	47
4.4 Conclusion: Strategic subject change and deception.....	48
5. Conclusion	51
6. Literature	52

“Language is not only the medium through which existence is transacted, it constitutes our central experiences of social and moral content, of such concepts as freedom and truth, and, most importantly, of individuality and the self; it is also a system of lies, evasions, propaganda, misrepresentation, and conformity.”

- Rachel Cusk, *Coventry: Essays*.

“No lie”

- Sean Paul ft. Dua Lipa – *No lie*

I. Introduction

In his *Principles of Philosophy*, which, incidentally, is the first philosophical text I have read during my studies, Descartes holds that “it’s not wise to place much trust in anyone or anything that has deceived us even once” (Descartes 2017, 1). Now whether this principle is a good one for the purposes of Epistemology is a question for a different master’s thesis. But for our everyday lives, it is *terrible* advice.

People deceive each other all the time. Children lie to not get punished for breaking household items, adults invent excuses for not showing up at their not-so-good-friend’s parties. Car manufacturers lie about the emissions their new models have. And when asked about our honest opinion about something, we have all found ourselves at least once carefully crafting an answer that might be true, but that most certainly makes our opinion seem more positive than it is. So, if we were to follow Descartes’ principle strictly, there would scarcely be anyone left we can trust.

Luckily, philosophical engagement with deception does not stop there. Philosophers have for a long time tried to figure out how exactly deception works, whether there is a relevant moral difference between different methods of deception, or what exactly it is that makes lying and the likes bad in the first place.

In recent years, there has been a sharp increase in philosophical interest in these topics. This has led both to the development of new arguments concerning existing questions, and to the emergence of new ones. The aim of this master’s thesis is to contribute to these recent developments, both by evaluating recent arguments as well as by bringing up new, relevant topics that have hitherto not received enough attention.

In chapter 2 of this thesis, I will begin by conducting a brief literature overview of recent developments in the philosophical work regarding deception. My focus here will be to sum up important work in the areas that I will later try to contribute to, specifically the linguistic and moral distinction between lying and misleading, and the emerging literature on plausible deniability.

In chapter 3, I will develop an argument concerning the moral status of lying and misleading. I will hold that in certain circumstances, misleading can be morally worse than lying, as it puts the victim of the deception in an epistemically weaker position to figure out that they have been received. This, as I will show, will often be morally relevant. Furthermore, I will evaluate several arguments according to which lying is generally morally worse than misleading and show why they are unsuccessful.

Finally, in chapter 4, I will explore an issue that has not yet been investigated by philosophers in a systematic way: speakers strategically changing the subject to a conversation. I will discuss what this strategy consists of and why it is frequently successful. Then, I will turn to attempts of speakers to avoid taking responsibility for norm-violating speech acts (including deception) that are currently discussed as cases of them using *plausible deniability*. I will argue that such cases are better understood as speakers trying to change the subject. Finally, I will turn to the topic of strategic subject-change in political speech and try to explain why it is such a common and frequently successful strategy in this domain. To end the chapter, I will discuss implications of this for the literature on deception.

2. Literature overview: Current developments in the philosophy of linguistic deception.

The last few decades have seen a significant increase in the philosophical engagement with deception. I believe there are two main reasons for this. The first concerns developments within philosophy. Important assumptions, such as the idea that lying necessarily involves an intention to deceive, or that lying is, for the most part, morally worse than misleading, have been called into question. This has both lead to new debates as well as novel contributions to existing ones. The other is more external. As the topic of misinformation and deception has taken center-stage in public discourse, many philosophers have turned to analyze its current manifestations. In this chapter, I will attempt to summarize recent developments in this research area.

I will focus my discussion on three main areas: The *linguistic devices* used for deception, their *moral status*, and *plausible deniability*, which is the most frequently discussed way in which speakers can avoid culpability for their deceptive speech. As we will see below, these areas cannot be neatly separated in practice, as results in one of them will influence theorizing in the others; nevertheless, I will use this distinction for purposes of readability, and discuss connections between the areas as they come up.

Before I get started, a brief terminological note: For the purposes of this literature overview, I will adopt a narrow understanding of what deception consists of; that is, I will be focusing on linguistic devices connected to communicating falsehoods; first and foremost, lying and misleading. Recent debates have led to theorizing concerning other rhetorical devices that some might describe as being deceptive in nature, such as dogwhistles/coded speech (see for example Khoo 2017; Saul 2024) or loaded questions (Davies 2016). My choice not to include these in this overview is purely practical, as there is already plenty to discuss without them. I will however touch on these topics in subsequent chapters (especially chapter 4), since a lot that can be discussed with regards to lying and misleading can also be applied to them.

2.1 Linguistic devices used for deception

Philosophers usually distinguish two main kinds of (linguistic) deception: Lying and misleading. This distinction was initially proposed as an ethical one (see below), but in recent years, its linguistic underpinnings have received more attention.

Historically, most philosophical investigations into deception have focused on lying. According to the received view, which can be traced back all the way to Augustine (1952), a speaker lies about a proposition *p* iff.

- (L1) the speaker says that *p*,
- (L2) they believe *p* to be false, and
- (L3) they intend to deceive their audience into believing *p* to be true. (Isenberg 1973, 248; Stokke 2013b, 348).

Before going into the main reasons philosophers have recently tended to abandon this definition, let me briefly touch on some earlier criticisms. The most common of these is that (L1) to (L3) are not jointly sufficient, as for a statement to be a lie, its content, *p*, also needs to be false (Coleman & Kay 1981; Moore 2000). However, most philosophers reject this proposal, as it seems to mainly be based on intuitions that most of them do not share (see Saul 2012b for more discussion).

Some have also argued that to lie about *p*, a speaker must also intend their audience to believe that the speaker believes *p* to be true (Frankfurt 1999; Simpson 1982). This too would mean that the above definition is not sufficient and would most likely suggest amending (L3). But in more recent definitions of lying, this has mostly been accommodated by changing (L1) into requiring that the speaker *asserts* *p*. Since assertion is usually taken to imply that the speaker in some way vouches for the truth of *p* (see Williamson 2000 for a more in-depth discussion), this is taken to imply that they communicate some degree of belief (or knowledge) that *p*.

Finally, it has been argued that (L3) is too narrow and thus not necessary for lying, as speakers need not intend to deceive their audiences about *p*; they can also intend for their audiences to falsely believe the speaker to believe that *p* (Chisholm and Feehan 1977). However, most philosophers hold that such cases should either be construed as misleading (see Saul 2012b), or that (L3), correctly understood, already captures such cases (see Williams 2002).

In more recent years, the most influential objection to the received view has been that an intention to deceive ((L3)) is not necessary for a speaker to lie. To this end, Carson (2006) and Sorensen (2007) present cases of what has become known as “bald-faced lies”, lies that clearly are uttered with no intention to deceive. One example of this, discussed by Carson, is that of a person who in front of a court lies about having witnessed a crime despite CCTV-footage clearly showing that they saw what happened, for fear of retaliation. In this case, the speaker cannot reasonably have an intention to deceive their audience, as there is well-known evidence that

defeats their testimony. Since however it still seems like they were lying, (L3) does not seem to be necessary for lying.

Most philosophers have accepted that such bald-faced lies are in fact lies and have attempted to provide new definitions that capture them¹. In what follows, I will discuss some of the most promising proposals.

Carson (2006) argues that one is lying only if one is “warranting the truth” of what one is saying. He argues that in the courtroom case, the speaker is still warranting for the truth of their testimony, since they do not qualify it in any way. But, as Fallis (2009) points out, this condition is too narrow. Intuitively, cases in which a speaker prefaces a lie by saying “I am not 100% sure, but” still seem like lies, despite the speaker not warranting the truth of what they are saying.

By contrast, Fallis (2009) argues that to lie, a speaker must assert *p*, which they believe to be false, while believing that they are in a context in which a norm of the form “Do only assert what you believe to be true” upholds. Similarly, Saul (2012b) takes lying to consist of saying something one believes to be false while in a “warranting context”, that is a context in which one is expected to be truthful. Here, a bald-face liar comes out as lying since, for example, a courtroom investigation clearly is a context in which it is well-known that one is to speak the truth.

But these accounts too have their problems. Stokke (2013a) argues that they falsely predict that most cases of irony will be lies, since they involve saying something that is literally false. I think this criticism can easily be extended to cases of speakers telling jokes, making suppositions of things they believe to be false (such as when presenting a *reductio-ad-absurdum*), engaging in pretense, or many cases of indirect speech. Perhaps, there are ways to spell out “warranting context” accounts in a way that excludes all circumstances in which any of these occur from the relevant set of contexts; but it is not clear how that would work, and it definitely needs further argument.

Finally, Stokke (2013a, 2016, 2018) has defended a view of lying according to which, very basically, to lie is to disrupt the goal of a conversation by means of assertion. But it is as of right now unclear whether this strategy is successful. Viebahn (2021) has pointed out that the underlying account of assertion is implausible when applied to non-lying cases; going into this however would require a much more detailed discussion which I am unable to provide here. And more generally, recent doubts that lying needs to involve assertion (see below) cast doubt on whether Stokke’s approach is a valid one.

¹ An exception to this is Lackey (2013), who argues that an intention to deceive that is broader than (L3) is necessary for lying.

Let us now turn to misleading. Most philosophers take misleading to be communicating something the speaker believes to be false by indirect means, most notably via conversational implicature. Viebahn (2021) discusses the helpful case of a doctor being asked about a patient that has recently died. He contrasts the following two possible (deceptive) answers:

(1) He is fine.

(2) The last time I saw him, he was fine.

(1) clearly seems to be an instance of lying. But (2) might be literally true – perhaps when the doctor last personally saw the patient, they were still fine. (2) then is a case of misleading: the speaker does not literally state or assert something they believe to be false, but given the context, they still implicate the false proposition that the patient is well.

Compared to lying, there have scarcely been attempts to give a unified definition of misleading (for a notable exception, see Stokke (2016, 2018)). Rather, misleading is usually treated as negatively defined as deception that is not lying. Additionally, while in principle, there is no reason why issues with lying should not similarly apply to misleading (for example, I do not see why there should not also be cases of bald-faced misleading), to my knowledge, those have not yet been systematically explored.

Still, some recent developments in the literature on lying have pressured philosophers into developing more substantial accounts of misleading. This is because some philosophers have contested condition (L1) of the received view of lying by arguing that speech-acts other than saying (or assertion) can be used to lie.² Holton (2019) and Keller (forthcoming) both argue that one can lie, rather than mislead, by presupposing something that is false. Viebahn (2020, 2021) likewise argues that not only can presuppositions be used to lie, but also questions, and even implicatures. The latter marks the biggest departure from received views, as implicature was the one kind of speech act where it was relatively uncontroversial that it could not be used for lying (cf. Berstler 2019). If such proposals are correct, it also has strong implications for views of misleading, as it means that misleading cannot merely consist in deceiving an audience by means other than lying.

These doubts about distinguishing lying and misleading merely by linguistic factors have led to views that mark a more radical departure from the received view of lying and misleading. The most notable of these is presented by Viebahn (2021) and Timmermann & Viebahn (2021). They

² Simpson (1992) defends a view according to which *every* linguistic device can be used for lying. Since this however leaves little room for an account of misleading, his theory has not proven as influential as these more recent proposals.

argue that what differentiates lying from misleading is the commitment a speaker makes to the proposition they intend to deceive their audiences about. Speakers are lying if they are strongly committed to the proposition, and misleading if they are weakly so. They argue that while in most cases, assertion comes with a strong commitment and indirect speech with a comparatively weaker one, this need not always be the case; hence, we should and can allow for cases where speakers lie via implicature or mislead via assertion.³

2.2 The moral status of deception

If we are to inquire about the moral status of deception, there are at least two kinds of questions we can ask. First, we can discuss just how morally reprehensible deceptive speech acts are, and why so. And second, we can discuss whether the different kinds of deceptive speech acts are morally on par, or whether there is a relevant difference between the two. Here, I will be focusing on the second of these questions, both for reasons of space and because the absolute moral status of deception seems to be much more disjointed from the debates discussed in the previous section.

There is a long-standing tradition in the ethics of speech according to which lying is morally worse than misleading. This view can be traced back to the work of Immanuel Kant. Kant, famously, argued that lying was morally reprehensible; in fact, so much so that it would even constitute a moral wrong if it was done to save someone else's life (Kant 1949). But in a passage in his *Lectures on Ethics* he discusses a case that while also morally questionable is, according to him, less problematic. The example he uses is that of a man packing a suitcase in front of another to make the latter believe that he was going away, while actually not planning to do so. Kant argues that this kind of deception is not as bad as flat-out lying about going on a journey. This is, he writes, because in the case of merely being misled, the inferences the audience draws are not the responsibility of the speaker (Kant 1963, 224f.).

Kant's exact argument to this conclusion is rather opaque and has been interpreted in at least two distinct ways. On the first interpretation, which in more recent years has been defended by MacIntyre (1994), there is a difference in duties governing lying and misleading. While there is a strong moral duty not to speak untruthfully, there is no corresponding duty against merely misleading someone, as inferences drawn from what a speaker says are not governed by any duties applying to the speaker.

³ For criticism of this view, see Marsili & Löhr (2022) and Pepp (2022).

While this interpretation is probably closer to what Kant believed than the one I will discuss in more detail below, it has not garnered nearly as much attention from philosophers. The reason for this is that it is not really an argument about lying and misleading. Rather, it merely interprets the two from the perspective of a prior moral theory. As such, it has no pull for those not subscribing to the that theory.

A more influential interpretation of Kant's argument is first defended by Adler (1997). According to him, what sets apart a lie from an act of misleading is that in the latter case, the hearer is partially responsible for being deceived. They are "[...] a rational, autonomous being and so fully responsible for the inferences [they draw]" (Adler 1997, 444). Since lying, or so it is assumed, does not require the kind of inferences misleading does, there, the responsibility is not split between speaker and hearer. Fricker (2012, 89) argues in a similar spirit:

If I tell someone that P, I personally offer my word to her that P, I take responsibility for P's truth. All she has to do, to come to believe P, is to take me at my word - accept what I offer. In contrast, if I tell her some other fact Q which [...] conversationally implies P, then if, in addition to trusting my word about Q, she also infers that P, that is, as it were, up to her. She, not me, is responsible for what I tell her - at least, she takes on at least half the responsibility for her coming to have the belief that P

What sets the Adler/Fricker interpretation apart from the MacIntyre one is that it depends less on a particular moral theory. All that the argument assumes is that bringing about a state of deception is bad, and that in cases of misleading, the responsibility for bringing about this state is divided between speaker and hearer.

Still, the argument has come under fire in recent years, and, to my knowledge, it has not been defended since. Most notably, Jennifer Saul (2012a, 2012b) argues that in most cases of misleading, the hearer does exactly what they are supposed to do when acting cooperatively. Take (2): It is clearly the doctor's intention to communicate that the patient is alive. Ignoring this would be heavily uncooperative and just counterintuitive. So, if a misled hearer does exactly what they are supposed to do so, and if they do so because the speaker wants them to, why should we assign them moral responsibility for the deception? A similar critique has also been advanced by Berstler (2019). Since I will discuss the issue more in chapter 3 and add objections of my own, I will not say more about this here.

The widespread rejection of the Kantian argument has led to new positions regarding the relation between lying and misleading. Saul (2012a, 2012b) argues that lying and misleading are generally morally on par, while allowing that in certain circumstances, one might be preferable to the other. Likewise, Viebahn (2021) and Timmermann & Viebahn (2021) hold that depending on

the wider context, lying or misleading might be the morally better option. It is however unclear to what extent their proposal differs from other positions, as they use a different account of the lying-misleading-distinction than most philosophers (see above).

Meanwhile, Berstler (2019), after rejecting the Kantian arguments, presents a novel argument for lying being worse than misleading. Drawing from Lewis (1975), she argues that assertion is governed by conventions of 'Truthfulness and Trust'. Furthermore, she holds that since these conventions are mutually beneficial and just, not following them is morally objectionable in and of itself. From this, she concludes that lying is worse than misleading, as lies violate the conventions of truthfulness and trust; meanwhile, she takes it that misleading is not governed by similar conventions and thus does not have the same moral status.

Viebahn (2022) objects to Berstler by arguing that not all lies are assertions; therefore, even if her general argument is correct, it fails to show a *general* difference in moral status between lying and misleading. However, I am not sure whether most philosophers will accept this criticism. While there have been arguments to include non-assertoric speech acts under the concept of lying (see above), most philosophers still hold that lying is fundamentally assertoric. In section 3, I will present my own objections to Berstler's argument, which do not depend on whether questions, presuppositions, etc. can be used to lie.

Of course, what I have discussed so far is only a subsection of the arguments regarding the moral status of misleading, namely, the ones that are discussed the most in the current literature. To name some other, less prominent arguments, Chisholm & Feehan (1977) hold that lying involves a breach of trust that misleading doesn't; but this is doubtful for the same reason Saul (2012a, 2012b) gives against the Kantian argument: Generally speaking, it seems that we also trust indirect communications, and so further evidence is needed to establish that lying violates that trust in a stronger sense.⁴ Webber (2013) argues that a liar loses all credibility, while a misleader might still be truthful in literal communication; Baumann (2015), in a similar vein, holds that misleading can still communicate basic respect for the hearer. But for both of these arguments, it is unclear whether they are really about lying and misleading themselves, or just about the moral character of the speakers using them (see Saul 2012b).

So far, most positions I have discussed here either hold that lying and misleading are morally on par and only differ in specific situations, or that lying is worse than misleading. A notable outlier from this is Rees (2014), who holds that misleading is worse than lying, since in misleading, I am

⁴ Strudler (2010) presents a similar argument, according to which more trust is placed in assertions, and thus lying comes with a bigger violation of trust.

using my audience in a morally objectionable way. I am relying on them to make an inference that leads to their deception. I will discuss this argument in more detail in chapter 3.

To summarize then, in recent debates, the position that lying and misleading are morally on par has become far more widespread. And even for those holding that misleading is preferable to lying, the grounds for holding this position have changed.

2.3 Deception and Deniability

In the final section of this brief literature review, I will turn to a topic that is closely connected to deception: (plausible) deniability. Plausible deniability, very broadly construed, is a speaker's ability to deny having meant some content *p* with an utterance *u*. Consider the following example by Pinker & Lee (2010, 785):

*Speeding Driver*⁵: A driver stopped for speeding says to the police officer:

(3a) I'm in a bit of a hurry. Is there any way we can settle this right now?"

When accused of offering a bribe to the police officer, the driver responds with:

(3b) I didn't mean to offer a bribe! I was just wondering if I could pay my fine right away.

In communicating his offer indirectly in (3a), the driver has maintained some deniability about trying to bribe the officer. He then makes use of this deniability in (3b), by claiming that he never intended to communicate the thing the officer accuses him of.

In this section, I will give an overview over debates concerning plausible deniability. First, I will discuss how plausible deniability and deception are connected. Then, I will touch on the main topic of contention among philosophers concerning deniability, namely what it is that gives speakers plausible deniability. Finally, I will turn to a problem regarding the scope of theories of plausible deniability.

While discussions surrounding plausible deniability are not limited to deniability about having deceived someone, a lot of the literature focussed on such cases. Walton (1996) discusses arguments where speakers make false claims, and, when called out for them, defend themselves by saying that they are merely quoting other people, or that they were merely making assumptions. Both Fricker (2012) and Peet (2015) discuss how having deniability about testimony might infringe on the knowledge granted by that testimony. And Saul (2024) points out how

⁵ I am here using Dinges & Zakkou's (2023, 374) reconstruction of this case, since it is a little neater than Lee & Pinker's version.

speakers who want to make false statements often frame them in such a way that they retain deniability about the falsehoods expressed.

There, I believe, are at least two reasons why it makes sense to discuss deniability with regards to deception. The first is that plausible deniability can help explain why speakers frequently get away with deception. Deceptive speech acts are almost always norm-violating in a very blatant way; after all, they straightforwardly deceive audiences. Yet it seems that despite the ubiquity of deception, the actual consequences speakers face for it are often quite limited. Plausible deniability, or so it is argued (cf. Khoo 2021), can explain why that is. When called out for their norm-violating speech, speakers can just claim not to have committed a norm violation or, at the very least, not to have intended it.

The second reason is that deniability might help us explain why speakers often chose to mislead rather than lie. If we consider these two options in isolation, misleading seems to be a worse choice: It is cognitively much more demanding than to just flat-out lie, and it can often be seen as impolite, or dodgy (Pinker & Lee 2010). However, since speakers usually have deniability about indirectly communicated content, and since misleading works by indirectly communicating something (for differing views, consider section 2.1), misleading rather than lying grants them deniability. This, in turn, might offset some of the disadvantages of misleading (see Pinker & Lee 2010; Camp 2018).

Let us now turn to actual debates about plausible deniability in the current literature. First, there is the question of what exactly it is that gives speakers plausible deniability. I will sort the views currently discussed on this into three categories: commitment-based views, epistemic views, and context-based views.

The first of these, commitment-based views, hold that a speaker has some degree of plausible deniability about some content p if they do not fully commit themselves to p . Walton (1996) is the most notable proponent of such a view. He argues that speakers have plausible deniability about propositions they merely suggest or bring forward by quoting others, since they do not commit to them in the same way they would if they asserted them. A similar view can be found in Fricker (2012), who argues that by implicating something, speakers do not commit themselves to it fully; she takes this to explain the deniability speakers have about implicated content. Given the recent emergence of views articulating the lying/misleading distinction in terms of degrees

of commitment (see above), we should expect this position to receive more attention in the future.⁶

Epistemic views of plausible deniability hold that what gives speakers plausible deniability about some content *p* is the epistemic relation between their audience and them having meant that *p*. Peet (2015) first articulates such a view in terms of a speaker being able to deny having meant *p* “with some degree of believability”⁷. And Dinges & Zakkou (2023) argue that speakers have plausible deniability about some content *p* iff. if they deny having meant *p*, the audience does not know that they meant *p*.

Finally, according to context-based views of plausible deniability, whether a speaker has deniability about having communicated some content *p* depends on the status their having communicated *p* within the conversational context. According to Camp (2018), a speaker has deniability about *p* iff. them denying having meant *p* is compatible with the propositions mutually acknowledged by the interlocutors at the time of the denial. Peet (forthcoming) proposes that for every set of accepted propositions within a conversation, there is a set of contexts compatible with those. If the speaker not having meant *p* is an option in a subset of the set of possible contexts, they have plausible deniability.

Let us now turn to a prominent problem that all these accounts face: That of adequate scope. This problem is generated by what Berstler (2019) calls “implausible deniability”. Implausible deniability describes cases where speakers seem to have *some* degree of deniability, but calling that degree “plausible” would be quite a stretch. Consider the example about the doctor asked about the wellbeing of a patient that has died, discussed above. In the example, the doctor must choose between the following two answers:

(1) He is fine.

(2) The last time I saw him, he was fine.

Assuming that communicating some content indirectly provides a speaker with more plausible deniability than communicating it directly (as most theorists do), we should hold that, at the very least, (2) would provide the doctor with *more* plausible deniability than (1). But this seems quite unintuitive; within the context of the question, it is clear that in uttering (2), the doctor intends to communicate that the patient still is fine. Calling (2) an instance of a speaker having *plausible* deniability would itself be implausible.

⁶ Boogart, Jansen & van Leeuwen (2020) also suggest that a speaker has an easier time denying something they have been less-than-fully committed to; however, they do not explore this position in detail.

⁷ Peet has recently defended a view that is more context-based; see below.

If this problem was limited to some cases of indirect communication, we could probably ignore them altogether, and fix the meaning of “plausible deniability” such that they are excluded. However, much of the literature in which the concept of plausible deniability is used seems to largely consist of such implausible cases⁸. For example, in discussions surrounding dogwhistles/coded speech, it is usually argued that one of the reasons speakers use them is that it provides them with plausible deniability (see Khoo 2018, 2021; Saul 2018, 2019, 2024). Likewise, Saul (2012b) holds that one of the reasons speakers mislead rather than lie is the deniability this grants them (see above).

This leads to a dilemma: Either we limit the notion of plausible deniability to cases in which the denial is actually plausible, thereby making it explanatorily useless for many of the cases it is applied to. Or we adopt a very wide notion, thereby including cases that intuitively are anything but plausible, and losing the distinction between plausible and implausible cases.

To my knowledge, the only attempt explicitly targeted at avoiding this dilemma is made by Dinges & Zakkou (2023). They introduce the concept of *untouchability*, which is supposed to apply to many cases where speakers are taken to *implausibly* deny that they meant some content. They define a speaker as having untouchability about having communicated some p iff. “if S denies that she meant [p] then it is practically irrational for S’s audience to engage in [hostile/opposing] actions or emotions” (Dinges & Zakkou 2023, 393).

I am however not sure whether this proposal is successful. In short, it still seems too restrictive to capture most of the cases in which plausible deniability is taken to play an explanatory role, despite actual denials seeming quite implausible. As Camp (2018) describes, actual cases of denial often lead to a back and forth between the speaker denying an interpretation of what they said, and their interlocutors insisting on it. If in all of these cases, the speakers had untouchability regarding having meant the content they are criticized for, this would mean that engaging in such a back-and-forth would be irrational.

Let us take stock: There is both a lively debate about what exactly plausible deniability consists of, as well as which cases should fall under the notion. As I will present my own thoughts on the latter question in chapter 4, I will leave it at that for now.

⁸ See chapter 4 for further examples.

3. Lying, Misleading and Epistemic Advantage

3.1 Introduction

Consider the following case:

Carina is cheating on her girlfriend Peggy with their mutual acquaintance Max. They normally only ever run into Max at the bookshop he works at. Peggy, developing a hunch that something is going on, asks Carina whether she has seen Max lately. Carina is torn between two possible responses:

(4) No, I haven't seen him.

(5) I haven't been to the bookshop lately.

Given just the above, both (4) and (5) would be morally problematic utterances, as with both, Carina would attempt to deceive Peggy. However, philosophical orthodoxy would have us say something more about these cases: (4), which is an act of *lying*, is worse than (5), which is an act of (merely) *misleading*. Misleading, or so the received view goes, is preferable to lying, either in virtue of a difference in moral status, or because a misleading speaker is less morally responsible for the deception than a lying one.⁹

This received view has been called into question, with some philosophers arguing that lying and misleading are morally on a par (most notably Saul 2012b). In this paper, my aim is to add to this critique. I will argue that a misled individual is in a worse epistemic position than one lied to regarding their deception. I will then show that this weaker position is sometimes morally significant. This suggests that there are cases in which misleading is in fact morally worse than lying.

The paper is structured in the following way: First, I discuss the linguistic differences between lying and misleading. I then present and criticize some recent arguments for the position that lying is morally worse than misleading. I go on to discuss a recent argument for misleading being worse than lying. Finally, I present my own argument regarding the epistemic status of the victims of misleading and discuss its moral consequences.

⁹ In the literature, these two related but distinct features are often not clearly distinguished and subsumed as lying being “worse than” misleading.

3.2 The lying-misleading distinction

If there is a moral difference between lying and misleading, this difference must supervene on linguistic differences. To see what these differences might be, let's first look at some more contrastive cases of lying and misleading. Consider Viebahn's (2021, 289) case of a doctor not wanting to tell a dying woman that her son has died (repeated from chapter 2.1). The doctor has to choose between the following two utterances:

(1) He is fine.

(2) I saw him yesterday and he was fine.

Or consider Bill Clinton's misleading answer when asked about whether he ever had an improper relationship with his former intern Monika Lewinsky. Saul (2012b, 89f.) contrasts his actual misleading statement in (6) with a fictional statement that would have been a clear case of lying in (5):

(6) There was no improper relationship.

(7) There is no improper relationship.

The most obvious difference between lying and misleading here is how the deceptive content is communicated¹⁰. In cases of lying ((1), (4) and (6)), speakers straightforwardly put forward a statement that they believe to be false. Meanwhile, in cases of misleading ((2), (5) and (7)), the speakers say something they believe to be true but that strongly suggests some additional content they believe to be false. Sam Berstler (2019) has attempted to capture this difference by arguing that when lying, a speaker asserts that *p* while not believing that *p*. When misleading, they conversationally implicate that *p* while not believing that *p*.

For the purposes of this paper, I suggest widening her account of misleading. On my proposal, misleading consists in *indirectly communicating* that *p* while believing that *-p*. The paradigm way of indirectly communicating something is via conversational implicature, but it also allows for other pragmatic devices to be means of misleading. This has two advantages: First, it captures cases discussed by Saul (2012b, 59f.) in which speakers produce utterances that by themselves are insufficient to communicate a definite truth-evaluable content to mislead their interlocutors.

¹⁰ Some philosophers have discussed cases of speakers lying with no intention to deceive (see e.g. Carson 2006 and Sorensen 2007). Since the purpose of this paper is to assess the moral features of lies *as* means to deceive, I will not consider such so-called "bald-faced lies".

Second, it is compatible with Stokke's (2016) position that both pragmatic presuppositions¹¹ and so-called prosodic focus too can be used to mislead.

To summarize, for the purposes of this paper I propose to conceive of lying as asserting that *p* while believing that $\neg p$, and of misleading as indirectly communicating that *p* while believing that $\neg p$. These are intended as minimal conditions for lying and misleading, respectively.¹²

3.3 Arguments for lying being worse than misleading

Again, if there is a moral difference between lying and misleading, then this moral difference must be grounded in linguistic differences between the two. But why should these linguistic differences make for any moral difference in the first place? Let's first look at some arguments that aim to establish that lying is worse than misleading.

3.3.1 The Kantian argument

The traditional view that lying is worse than misleading is commonly traced back to Kant. In his *Lectures on Ethics*, he introduces the case of a man packing his suitcase in front of another to make the other man falsely believe that he is going travelling. Kant then argues that while deceptive, this kind of behaviour is not morally questionable in the same way lying is. The audience has no right to expect that the inference drawn from what they see is true¹³.

Philosophers have frequently applied this idea to the lying-misleading-distinction: Since in being misled, the hearer is required to make an inference to the deceptive content, they are partially responsible for their deception. This then mitigates the responsibility of the speaker (see for example Adler 1997).

Elizabeth Fricker (2012) has put forward an argument in the same spirit:

If I tell someone that *P*, I personally offer my word to her that *P*, I take responsibility for *P*'s truth. All she has to do, to come to believe *P*, is to take me at my word - accept what I offer. In contrast, if I tell her some other fact *Q* which, in our mutual *K*-context¹⁴, conversationally implies *P*, then if, in addition to trusting my word about *Q*, she also infers that *P*, that is, as it were, up to her. She, not me, is responsible for what she infers from what I tell her - at least, she takes on at least half the responsibility for her coming to have the belief that *P*. (Fricker 2012, 89f.)

¹¹ For a different position, according to which one can *lie* by presupposing something one believes to be false, see Holton (2019), Keller (forthcoming) and Viebahn (2021).

¹² For the broader debates about what exactly lying and misleading consists of which I will not consider here, see chapter 2.

¹³ Found in the L. Infield translation (1930), 226f.

¹⁴ "K-context" is short for "knowledge-context".

To assess this argument, we first need to get clear on how “responsibility” is used here. Arguments in the Kantian tradition frequently rely on a mislead individual being partially responsible for their deception in some intuitive, *non-moral* sense. The same is not true for an individual lied to. They then infer a difference in moral responsibility from this difference in non-moral responsibility. In order not to mix up the various responsibilities at play, I will refer to the latter as “agential responsibility”, as it is usually connected to an individual playing an active role in bringing about their deception.

Applying this distinction, we arrive at the following reconstruction of Fricker’s argument:

P1: In cases of lying, the speaker is fully agentially responsible for the audience’s deception.

P2: In cases of misleading, the audience is partially agentially responsible for their own deception.

P3: If a speaker is only partially agentially responsible for the audiences deception, they are only partially morally responsible for it.

C: A misleading speaker is only partially morally responsible for the deception.

A standard way of critiquing Kantian arguments of this sort is to dispute P3. For example, Saul (2012a, 4f.) argues that in general, this relation between agential and moral responsibility does not hold. In many cases, the victim being partially responsible for the misdeed committed against them does not reduce the perpetrator’s moral responsibility. However, showing that P3 does not *always* hold is insufficient to debunk the argument, as it only requires that the relation holds for cases of *deception*.

Thus, to see whether Fricker’s argument is sound, we need to investigate *how* acts of misleading generate shared agential responsibility (as in P2) and only then whether that explanation justifies the claim that the shared agential responsibility leads to shared moral responsibility (as in P3).

Berstler (2019, 12f.) reads Fricker as arguing that victims of misleading are partially agentially responsible for their deception because they must perform more cognitive work to be deceived than victims of lying. This is a rather uncharitable reconstruction, since the underlying assumption is often false: victims of a very confusingly phrased lie need to perform more cognitive work than those of a very straightforward act of misleading. Furthermore, this reconstruction is at odds with Fricker saying that in cases of lying, the perpetrator (wholly) “takes responsibility for [the purported] truth”. If the difference were one of degree, not *all* the agential responsibility could lie with the speaker when lying.

A more plausible way to reconstruct Fricker is that she supposes that there is a difference in *kind* between the inference required to arrive at the misleading content and the inference required to arrive at the content of a lie. She writes that a hearer's inference in cases of misleading is "up to her". I see two ways to interpret this phrase. Either Fricker claims that the relevant inference is *optional*: It is something that is performed in addition to interpreting the literal content of their interlocutor's statement. Since the hearer did not *have to perform* this additional step of interpreting, she is (partially) responsible for it occurring. Or she holds that for indirectly communicated content, it is more *indeterminate* what that content is, when compared to literal content. The hearer is partially responsible for their deception because they interpreted the speaker to mean something *particular* rather than something else they could have taken them to mean.

However, neither of these options should make a moral difference, thus neither is sufficient to support P3. Let's begin with the *optionality* thesis. While it is technically true that hearers are not *required* to work out implicatures or other indirectly communicated propositions, this makes for a very unrealistic portrayal of communicative practices. Implicatures often arise precisely because on its own, the literal content of an utterance fails to be relevant or otherwise communicatively cooperative. For example, if

(5) I haven't been to the bookshop lately.

would *not* be taken to communicate that Carina has not recently met Max, she would not have answered Peggy's question; calculating the relevant implicature here is only natural. In fact, we should assume that in most cases of intentional misleading, the relevant communicative contents come very naturally to the hearer; otherwise, they would not be particularly clever acts of misleading. But if hearers are expected to work out the relevant implicatures, it is unclear why they should be blamed for them.

This question becomes all the more pressing if we compare cases of misleading to other cases of picking up indirectly communicated contents. If the fact that working such contents out is *optional* alters the interpreter's moral responsibility in deceptive cases, we should expect similar alterations in responsibility in cases where no deception is involved. But this does not hold up. Assume there are two separate folders that a form might go into. I ask you which one is the correct one, and you answer me with

(8) It does not go into the left one.

If I were later to claim that I figured out – by myself! – that the file goes in the right folder, you would rightfully be irritated. Clearly, I had no clue. But you did not *assert* that it goes in the right one, you merely implicated it. Still, despite the implicature, it would be wrong to say that I was (partially) responsible, and thus praiseworthy for my success. To summarize, not only is the *optionality* thesis insufficiently unmotivated; it gives us the wrong kinds of results in other cases.

Let's turn to the *indeterminacy* thesis. According to it, the misled individual is partially responsible for their deception because they inferred a particular indirectly communicated content rather than a different one. They could have arrived at a different interpretation that did not entail their deception.

There are two problems with this idea. First, if the issue with implicatures is that the hearer could have worked out a different implicature than they did, we should expect similar results for certain kinds of lies. Assume a murder mystery has happened in an old Victorian mansion, and a detective is tasked with solving the mystery. Two women, call them Anna and Beth, are standing next to each other. An accomplice of Anna wants to deter the detective from figuring out that she is in fact the murderer. In a context where it is pretty obvious from contextual factors that the accomplice wants to talk about Anna, he says to the detective

(9) I know for a fact that she did not do it.

Clearly, the accomplice is lying to the detective. But if *indeterminacy* is correct, we should argue that the accomplice is not wholly morally responsible for the deception. After all, the detective could have interpreted (9) as being about Beth. So, if *indeterminacy* is correct, it is not clear why it be applied only to cases of misleading.

Second, it is again unclear why *indeterminacy* should make any difference with regards to moral responsibility in the first place. Even if a misled individual could have arrived at a different interpretation of their interlocutor's utterance, that does not mean that they should have done so. After all, in the case of being intentionally misled, they do exactly what the speaker wants them to. They interpret their speech in exactly the way they plan them to.

Similarly to *optionality*, *indeterminacy* does not just seem to be ill-justified, it also gives the wrong results for parallel cases. To use the example of the files and (8) again: Of course, I could have interpreted your contribution as saying something else; but it would be wrong to claim that the fact that I did not do so makes me partially responsible for the correctness of my filing.

To summarise then, even if Fricker is technically correct that there is a difference between how assertions and acts of indirect communication are interpreted, that does not mean that this

technical difference amounts to a difference in moral responsibility. Thus, P3 seems at best to be unmotivated, but given the problems outlined above, is more likely false.

3.3.2 Berstler: Lying is unfair in the face of conventions

Sam Berstler has proposed a novel argument for the conclusion that lying is in fact worse than misleading, based on Lewis' theory of conventions. According to Lewis' account, the meaning of our expressions is explained via two conventions of *Truthfulness* and *Trust*, which speakers of a language for the most part conform to. Berstler (2019, 18) spells these conventions out as follows:

Let *x* be a sentence and *p* a proposition. If *x* (conventionally) means *p*, then:

Truthfulness: (a) Speakers assert *x* only if they believe *p*.

Trust: (b) If speakers assert *x*, their addressees come to believe *p* (unless they have some other reason not to).

She then goes on to introduce a Nozickian notion of fair play, which is meant to spell out the kind of moral obligations involved in participating in certain kinds of cooperative ventures:

Fair Play [...]: When a number of people engage in a just, mutually advantageous cooperative venture, according to rules, and thus restrain their liberty in ways necessary to yield advantage for all, and when you have benefitted from these people's submission, you are obligated to conform to the rules. (Berstler 2019, 21)

Berstler's argument, then, is that lying is worse than misleading because a liar makes use of the mutually advantageous convention of assertion¹⁵ while not conforming to its rules (Berstler 2019, 17f.). They are playing the assertion game unfairly. The same, cannot be said about misleaders: Since the devices used in misleading are not conventional, no rules of any sort are violated here. This, she takes it, leads to a difference in moral status between lying and misleading. A liar violates additional obligations; thus, they commit a graver wrong.

The problem with this argument, I think, is that Berstler does not give us good reasons to believe no fair play obligations are violated in misleading. She argues that this is so because there are no metasemantic conventions governing the indirect speech acts involved in misleading. However, this justification is insufficient to rule out that they meet the requirements for Fair Play. Berstler

¹⁵ It might be that Berstler gets a bit lost in the ambiguity involved in the term "assertion" here. Lewis proposes his conventions of Truthfulness and Trust as conventions about *declarative sentences*, not about some specific illocutionary act they are used in. Meanwhile, lies clearly are about the illocutionary level. For reasons of space, I will not discuss this issue here further; but overall, we should be doubtful of the connection between metasemantic conventions and illocutionary acts Berstler uses in her argument.

takes it that Fair Play obligations apply to practices in case they are *cooperative*, *beneficial*, *just*, and in case participation in them is *voluntary* (Berstler 2019, 21f.). Clearly, these requirements can be met by all sorts of practices, not just those governed by Lewisian conventions. Thus, Berstler's argument leaves open the possibility of misleading violating the same kinds of obligations as lying. It fails to prove a difference in moral status between the two.

In fact, I would argue that the speech acts involved in misleading *are* subject to Fair Play obligations. There is a practice of speakers only intentionally indirectly communicating what they believe to be true. If this were not the case, misleading would simply not work, as relying on the indirect content of your interlocutor's statements would be an arbitrary choice.

Furthermore, this kind of practice meets Berstler's requirements for having Fair Play duties attached to them. It is *cooperative*, as its being maintained depends on a wide range of speakers and hearers acting according to it. It is *beneficial*, as it allows speakers to use a wider range of communicative devices, and it enables hearers to trust and act in accordance with indirectly communicated contents. It also makes communication more efficient, as we do not have to state every little piece of information explicitly. It is *just*, as it does not discriminate between who the speaker or audience is. And it is *voluntary*. Save for certain legal contexts, I am not forced to implicate only what I believe to be true.

Thus, judging by Berstler's own principles, the indirect speech acts involved in misleading are subject to Fair Pay obligations. These obligations are violated whenever a speaker intentionally misleads. Therefore, there is no difference in moral status between lying and misleading based on Fair Play obligations. Both are similarly guilty of being unfair.

3.4 An argument for misleading being worse than lying

Let's now turn to an argument for misleading being worse than lying, introduced by Clea F. Rees. She argues that to believe and therefore be deceived by a liar, their interlocutor only needs to assume *Truthfulness* and *Trust à la Lewis* (see above). Meanwhile, to interpret misleading utterances as the speaker intends, they also need to assume Gricean communicative cooperation. From this, Rees infers that a misleader relies on the misled individual's reasoning in a stronger way than a liar does and, thus, *uses* their interlocutors in a *more* morally problematic way:

It is in virtue of her own moral and epistemic goodwill that the merely-deliberately-misled presumes the goodwill of fellow conversants which licenses her assumption that her deceiver is cooperating with her in furthering shared conversational ends, that her deceiver's assertions are therefore intended to further these ends, and that her deceiver's false conversational implicatures are thus as trustworthy as her assertions. Far from

mitigating her deceiver's culpability, therefore, the greater and more active responsibility of the merely-deliberately-misled for her own deception actually exacerbates the wrongness of deceiving her. The mere deliberate misleader uses the deceived in an epistemically and morally objectionable way that the liar does not. (Rees 2014, 62)

Here is a more formal way of looking at this argument:

P1: Individuals who are misled are made agentially more responsible for their deception than individuals who are lied to.

P2: Making an individual agentially responsible for their own deception is morally wrong.

C: Therefore, misleading is morally worse than lying.

Contrary to Fricker's argument, to which it is otherwise similar, Rees takes it that not only are misleaders more responsible for the deception, but that misleading is genuinely worse than lying, since in addition to deceiving them, using an individual is a wrong of itself.

P2 is a normative premise that within Rees' argument is taken for granted. When evaluating it, then, we need to discuss whether P1 is correct and, even if it is, whether the kind of responsibility necessary for P1 is the kind applicable in P2.

Rees herself supports P1 by arguing that in addition to being trusting, a victim of misleading (but not of lying) needs to be conversationally cooperative in the Gricean sense to arrive at the inference that the deceiver wants them to make. This justification, however, is mistaken, as the same is clearly true for victims of lying. Take the Gricean maxim of Quality, which Grice splits into two submaxims, "Do not say what you believe to be false" and "Do not say that for which you lack adequate evidence". Clearly, for a lie to be successful at deceiving, the audience must assume the speaker to be following these maxims, as otherwise, they would have no reason to believe the lie. Likewise, the audience must assume that the liar follows the maxim of Relation (later often "Relevance"), with the single maxim "Be relevant", for otherwise, there would be no reason why the audience would have to assume that the thing the liar said was relevant to them in any way (Grice 1975, 46).

It seems, then, that in order for Reese's argument to work, we need to find a different justification for P1. But matters are not quite so simple, as it is not sufficient to find *some* difference in agential responsibility between lying and misleading. The alleged difference must be of the kind to plausibly be an instance for which P2 is true. Say I am sending a friend an article with deceptive information. If I send them the link to the article rather than the article itself, I make my friend *more* agentially responsible for their deception. They have a marginally larger role in bringing it

about – they have to click the link themselves. But clearly, in this case, the fact that they play a larger role in their deception should not amount to a moral difference.

Given this caveat, I do not see how the differences in agential responsibility hitherto discussed are useful for supporting Rees's argument. Sure, misled individuals make an optional or indeterminate inference (see the section on Fricker's argument). But as discussed before, such inferences are so very commonplace when interpreting our interlocutor's conversational contributions, it would be highly implausible to attach the kind of normative significance required for the argument to them. It seems then that without some different justification for P1, Rees' argument fails.

3.5 The epistemic consequences of misleading

In this section, I will present my own argument for misleading differing from lying in ways that are often morally significant. The argument will proceed in two stages. First, I will show that misleading puts the victim in a weaker epistemic position than lying does. Second, I will argue that in many situations, this weaker epistemic position can be morally relevant.

To see how misleading has potentially negative epistemic consequences, let's return to the case of Carina and Peggy. Say Carina chooses to lie to Peggy about not having seen Max lately, and thus answers her with

(4) No, I haven't seen him.

If Peggy believes (4), Carina will have succeeded in deceiving her. But, as it is often the case, this need not be the end of the story. Assume Peggy, at a later point, learns that Carina in fact did meet Max. Not only will she now know that her previous belief was false, but she will also know that Carina *intentionally deceived* her. After all, whether Carina lied in asserting (4) solely depends on whether she believed the literal content of (4) to be false. Once Peggy learns that (4) is false, she thereby automatically knows that Carina lied, and thereby intentionally deceived her.

Compare this to the alternative case of Carina uttering

(5) I haven't been to the bookshop lately.

and thereby misleading Peggy. In this case, once Peggy learns the truth, she too will know that her previous belief was false. However, she will not automatically know that Carina intentionally deceived her. Of course, she will know that what she took (5) to mean (namely that Carina had not met Max lately) is false. But in order to then know that (5) was a case of *intentional* misleading rather than a mere misunderstanding, she also needs to know that this in fact *is* what Carina

meant by it. And knowing this requires a whole lot of additional information: Carina must remember their conversation before (5), and she must be aware of any contextual factors at the time that might have influenced what could be implicated by uttering (5). Remembering these details is not a given, especially if the original conversation happened a while back.

What this example shows is that there is an *epistemic difference* between victims of lying and victims of misleading. Misled individuals are in a weaker epistemic position to know that they have been intentionally deceived. This is because in being misled, the deceptive content was communicated indirectly. Therefore, it is context-dependent to a much larger degree than the literally asserted content of a lie. Consequently, misled individuals need to know that they are not mistaken about the context of an utterance, in order to know that it was not a case of intentional misleading.

This epistemic difference is a difference in degree. In many cases, figuring out that a speaker lied will be much more difficult than in the case of Carina and Peggy. But even in those cases, lied-to individuals are at a (comparative) epistemic advantage: They do not need to be aware of certain contextual features to figure out that they have been deceived. The same is not true for misled individuals.

Furthermore, the epistemic difference does not stop there. Being in a better or worse position to find out you have been deceived will plausibly carry over to other (partly) epistemic features of the situation. For example, provided they did not put in some extra work, a misled individual will be in a worse position to state that they have been deceived. They will need to provide more evidence to prove the wrongdoing of the other, or else make a less believable statement. Or consider actions they want to take because of being deceived, such as accusing the other of wrongdoing, altering their conduct towards them, etc. If the justification of such actions depends in any way on the epistemic status of the agent, a misled individual will be less justified than a lied-to individual.

So far, I have argued that there is an epistemic difference between being lied to and being misled. I will now argue that plausibly, this epistemic difference is often morally relevant. To illustrate this, I will describe three types of cases in which this relevance is particularly striking.

First, in many cases, individuals will non-instrumentally care about knowing whether they have been deceived. Take the case of Carina and Peggy again: It is very likely that they value honesty within their relationship. Therefore, Peggy will not just want not to be deceived, she will also want to *know* whether she has been deceived, and thereby, whether her relationship's commitment to honesty is being honoured. In such situations, it is very likely that the epistemic differences

discussed above have moral relevance: By making it harder to figure out the intentional deception, the misleading individual infringes on what their audience cares about.

Second, in some cases, putting the audience in a more difficult epistemic position is part of the speaker's intentions. Assume that Carina expects Peggy to eventually figure out that she has an affair. Perhaps Carina then might mislead rather than lie to Peggy precisely *because* of the epistemic status that would put Peggy in. Even if Peggy figures out that Carina was cheating on her, she might not immediately realise that Carina also deceived her about it if she does not remember the context of Carina's misleading utterance. Furthermore, even if she does, she will have a harder time accusing her of it, than if she had lied. In this case, then, the epistemic consequences of misleading would not merely be a byproduct of Carina misleading rather than lying; they are part of her intentions when speaking. And if achieving specific epistemic consequences is part of a speaker's intentions, it is very likely morally relevant.

Finally, several moral philosophers have argued that wrongdoing changes the normative landscape by creating new rights and duties. For example, Sliwa (2020, 79) argues that wrongdoing creates *reparative duties* toward the wronged person. These include apologizing to them, acknowledging the wrongdoing, and making amends. The wronged person meanwhile requires corresponding *reparative rights*. Bennett (2018) likewise argues that wrongdoing creates what he calls *secondary duties* for wrongdoer and, potentially, others. If this general idea is correct, it applies to cases of morally problematic deception, too: A deceiver should, for example, apologize, and a deceived individual has a right to have the wrong acknowledged.

The epistemic difference between being lied to and being misled also matters when it comes to such reparative duties and the corresponding rights. Consider the case of apologies. Assuming the deception was morally problematic, both a lied to and a misled individual have a right towards an apology. But the misled individual is in a much tougher position to actually receive it. As discussed above, it is more difficult for them even to notice that they have been intentionally deceived. Moreover, they might have troubles demanding an apology, as they cannot just point to what was said to make their right to one clear. They must also provide adequate context to what happened. Since the right towards an apology (or other reparative rights) is a *moral* right, difficulties in having it met will probably have some moral import.

Let's take stock: Misleading puts the victim in a worse epistemic position than lying. And the role of these epistemic consequences is such that they often are plausibly morally relevant. Now whether or not this moral relevance amounts to misleading being morally worse than lying in a given case most likely depends on other features of that case and of the moral theory used to

analyse it. But if we take together the fact that arguments for lying being worse than misleading are frequently mistaken, and the fact that misleading often has morally relevant, bad consequences, we have good reason to believe that misleading, at least in some cases, is morally worse than lying.

3.6 Conclusion

In this chapter, I have argued that misleading puts the victim in a worse epistemic position than lying. I have also shown that this feature of misleading is often morally relevant. Together, I take it, this gives us good reason to believe that misleading is, at least sometimes, morally worse than lying. I will now briefly discuss two consequences this argument has.

First, some philosophers have recently defended views according to which there is no general difference in moral status between lying and misleading. Rather, lying and misleading have different properties, and their respective moral status depends on the role these properties play in specific circumstances. For example, Timmermann & Viebahn (2021) describe the difference between lying and misleading as one of degrees of commitment. Furthermore, they argue that depending on the circumstances, being more or less committed to a proposition can be morally preferable.¹⁶ The argument I have presented in this paper both supports this general approach towards the lying and misleading and suggests a relevant further feature to investigate in this regard: The epistemic consequences of deceptive speech acts.

Second, recognising that misleading can be worse than lying because of its further effects on the moral landscape forces us to take a perspective that is currently largely absent from discussions in the ethics of language. While debates about what makes, say, slurs or acts of silencing bad are ubiquitous, such debates hardly ever consider what happens after the reprehensible speech act. Recognising that what is likely to follow from deception is morally relevant can serve as a gateway to investigating similar aspects of speech acts that have hitherto been neglected.

¹⁶ See chapter 2 for more a more detailed explanation of this position.

4. Strategic Subject-change, plausible deniability, and dodgy politicians

4.1 Introduction

In ordinary conversation as well as in more official contexts, speakers often try to avoid talking about certain subjects. Consider the following three cases:

Moving in: Carina and Peggy have been dating for several months. During a conversation about their shared future, Carina asks Peggy whether she would like to move in with her. Peggy does not feel ready to do so yet, but at the same time, she does not want to say no. She wants to avoid hurting Carina's feeling, and she thinks that she might change her opinion after giving it some more thought. She replies:

(10) I'm a bit short on money right now; I am not sure whether I could afford moving.

Avoiding job talk: During a night out with friends, Clark is talking to Nelly, whom he has a crush on. Eventually, she asks him about his job. Clark has recently been fired. He is embarrassed about this and does not want Nelly to think less of him. Turning to a friend of his, he asks:

(11) I just remembered, Martin, you recently got a dog, right?

Vaccines and parties: Following accusations of hosting a party in violation of covid laws at his flat, then UK prime minister Boris Johnson was asked by an MP what was going on at his flat in Downing Street on the alleged date, Johnson responded by saying:

(12) I can tell you what was going on at Downing Street. [...] The process must go on. We've been delivering the fastest vaccine and booster rollout in Europe. We've been getting people back into work. We've been helping to bring the West together.¹⁷

While these three cases differ in many ways, they all are instances of the same general strategy: A speaker, confronted with a question they do not want to answer, changes the subject.

Call this phenomenon *strategic subject-change*. In everyday discourse, strategic subject-change is ubiquitous. Most of us have at some point probably found ourselves in a similar situation as Peggy or Clark, and we are all too familiar with politicians dodging questions or criticism by

¹⁷ <https://www.thecourier.co.uk/fp/politics/scottish-politics/2970936/video-watch-boris-johnson-dodge-question-on-downing-street-party-allegation/> (last accessed 15.05.2024).

evading a direct answer. Yet, perhaps surprisingly, philosophers of language have hitherto not discussed strategic subject-change in a systematic way.

In this paper, I will begin to fill this gap. In the first half of this paper, I will identify some core characteristics of strategic subject-change. This will enable us to explain how strategic subject-change can be a successful strategy for avoiding unfavorable topics. In the second half of the paper, I will apply these findings: First, I will argue that understanding strategic subject-change can help us solve a puzzle in the literature on plausible deniability. Then, I will explore the issue of subject-change in political speech and give some reasons why it is especially common in this domain.

4.2 What is strategic subject-change?

We strategically change the subject all the time. We decide to talk about economics rather than migration policies once our racist uncle joins the conversation. We make an innocuous comment about something that is happening in the room we're in rather than talking about that one embarrassing incident our spouse keeps bringing up. Or we shoot back in an argument because it would hurt our pride to admit to the wrongdoing we are accused of.

In this section, I will first discuss some main characteristics of strategic subject-change. I will then turn to what often makes it a successful strategy for avoiding uncomfortable topics.

4.2.1 What is strategic subject-change?

In this section, I will develop an idea of what strategic subject-change is. My goal is not to provide a full definition, but rather to give an intuitive characterization that helps us understand why cases like the ones discussed above are successful, and why they are so ubiquitous.

Let us start with the obvious: strategic subject-change is an instance of changing the subject. When we engage in conversations, we are not randomly throwing out thoughts as they enter our heads. We focus on topics. We decide on plans, we try to answer a question, we exchange information about an incident. And while our conversations initially revolve around a specific subject, they do not necessarily stick to it. Speakers bring up new, often related topics; they focus their conversation on a part of the initial topic; or they feel like they have exhausted one topic and bring up another. The speech acts that lead to such changes of topic constitute subject-change.

One way in which instances of subject-change can differ is by the way in which the attempted subject-change is brought about. Let us distinguish two broad categories:

Direct subject-change: The speaker explicitly proposes the subject-change, by raising a new question or introducing a new topic.

Indirect subject-change: The speaker produces an utterance that presupposes a subject-change.

The difference should be familiar from everyday conversations. If I want to bring up the topic of my new car, I can do so directly, by saying something like “Can I tell you about my new car?”. Or I can do so indirectly, by just stating “I have a new car”, which presupposes that my new car is topical. Both of these bring about the same subject-change, they just do so by different means.

Let us now turn to characteristics of *strategic* subject-change. Intuitively, the main way that strategic uses of subject-change differ from non-strategic ones is the *goal* with which they are performed. In the cases discussed above, speakers change the subject because they want to *avoid* the current topic. They are driven not by a mere desire to introduce a new topic, but to move away from the old one. More succinctly:

Topic avoidance: Strategic subject-change is performed with the goal of avoiding the current topic of the conversation.

Importantly, *Topic avoidance* does not state that avoiding a topic is the *main* or *only* goal with which strategic subject-change is performed. Consider *Moving In*. Peggy’s main goal is to avoid hurting Carina’s feelings. She aims at avoiding the topic of moving in because doing so is part of this larger goal. Likewise, in *Vaccines and parties*, changing the subject is not Johnson’s *only* goal. He also, presumably, wants to use his statement to convince the members of parliament listening that he is doing a good job as a prime minister.

Still, *Topic avoidance* is strong enough to distinguish such strategic uses from ordinary subject-change. Take the case of a group of friends discussing their upcoming trip to the zoo, and then switching the topic to one of the members’ annoying father-in-law. Here, the speaker bringing up the topic of his family wants to change the subject away from the zoo visit, but not because he wants to avoid that topic. Rather, he thinks that the topic is done being discussed, and he has an interest in bringing up a different one. But this interest has nothing to do with the previous topic.

But *Topic avoidance* by itself is not sufficient to differentiate strategic from mere subject-change. Consider an alternative version of *Moving In*, in which Peggy responds to Carina’s question by saying:

(10*) I don’t feel ready to have this conversation. And besides, I don’t really have the money to move right now.

Here too, Peggy changes the subject. She also *Topic avoidance*, as she changes the subject because she wants to avoid Carina's question. But intuitively, as she openly admits that she wants to avoid it, it does not seem strategic in the same way as the other cases. This suggests that we should add a further characteristic of strategic subject-change:

Obfuscation: Strategic subject-change is performed in such a way that it is difficult for hearers too identify its goal of avoiding the current topic of the conversation.

Together, *Topic avoidance* and *Obfuscation* give us a decent understanding of what strategic subject-change consists of. In the next section, I will use this understanding to show how it can be a successful strategy to avoid undesirable topics.

4.2.2 Why is strategic subject-change successful?

Strategic subject-change is ubiquitous, both in everyday conversations and in high-stakes contexts such as interviews, political discussions and hearings. Thus, we should assume that it is frequently successful. In this section, I will try to explain why it is. To do so, I will first discuss what it means for strategic subject-change to succeed. Then, I will explore several ways in which it can do so.

4.2.2.1 What does it mean for strategic subject-change to be successful?

In the previous section, I have argued that the constitutive goal of strategic subject-change is to avoid the current topic of a conversation. In this section, I will spell this out in more detail; to do so, I will make use of Roberts' (2012) influential "question-under-discussion" framework.

As we have already seen in the previous sections, conversations revolve around topics. Following Roberts (2012), we can understand such topics as questions-under discussion (QUDs), that is, questions the participants are jointly trying to answer. So, for example, when deliberating where to go for dinner after a workshop, participants are jointly trying to answer the QUD "Where should we go for dinner?".

If a new topic is raised, a corresponding QUD is added to the conversational context. There are two ways in which this can be achieved. First, a question can be raised directly. This matches what I have called *direct subject-change* above: A speaker either asks a question or introduces a new topic explicitly. In the case of our hungry philosophers, this could happen by one of the participants asking "Where should we go for dinner?", or by asserting something like "We should decide on a place to go eat".

Second, a QUD can be added by being *accommodated*. If someone makes a contribution to the conversation that presupposes a certain question under discussion, and if that contribution is accepted, the presupposed question will be added to the conversational context (Roberts 2012;

see also Lewis 1979 and Stalnaker 2002 conversational accommodation more generally). For example, rather than asking where they should go for dinner, a participant of the workshop could just suggest “We should go to [Add punny restaurant name] after this!”, and if her interlocutors accept that contribution to the conversation, the corresponding QUD will be added to the context. This matches what I called *indirect subject-change* above.

Roberts also identifies two ways in which questions can be removed from the conversational context: either, they are answered by the participants, or they are deemed unanswerable. But for the purpose of describing the way actual conversations proceed, these options do not seem exhaustive. Sometimes, speakers will drop a question under discussion because another one seems more pressing or salient. And sometimes, they will simply move away from a topic without even noticing that the QUD has changed. We are all too familiar with the kind of meandering conversations that lead from one topic to another, without any explicit topic change ever being acknowledged.

Now that we have a firmer grip on how conversational subjects work, let us return to strategic subject-change. As I have pointed out above, strategic subject-change is aimed at avoiding the current topic of the conversation (*Topic avoidance*). In terms of the QUD-framework, this means that a speaker engaging in strategic subject-change makes an attempt to remove a question they do not want to discuss from the conversational context.

But since speakers engaging in strategic subject-change both want to avoid the current topic (*Topic avoidance*) and want this motivation to be hidden from their audience (*Obfuscation*), not just any way of removing the corresponding question from the conversational context will do. Answering it would mean having to engage with it – this is exactly what speakers using strategic-subject change want to avoid. And deeming it unanswerable would require them to disclose that they do not want to answer it, which is at odds with *Obfuscation*.

So, for strategic subject-change to be successful, speakers need to replace the current QUD with a different one, without answering the former or deeming it unanswerable. In the next sections, I will explore some central ways in which this can succeed.

4.2.2.2 The subject-change is accepted

The utterances that constitute strategic subject-change are always at odds with the current QUD. Take *Avoiding Job talk*. By asking Martin about his new flat, Clark is attempting to add a new QUD to the common ground without first answering the current one. This can also happen indirectly: By uttering (10), Peggy is making an utterance that does not appear relevant with regard to Carina’s precise question. Thereby, she attempts to indirectly change the QUD.

The most straightforward way in which such attempts at changing the subject can succeed is for them just to be accepted by the interlocutors. They go along with the directly suggested topic, or they accommodate the indirectly communicated one.

There are several reasons why this might happen. First, interlocutors might not even notice the subject-change. As I have mentioned before, in many everyday conversations, it is perfectly natural to switch topics without acknowledging doing so. In such a situation, strategic subject-change might be perfectly fine with the participants.

Second, it is not always clear what the current question-under-discussion is. Take *Moving In*. While we can interpret Carina's as leading to the QUD "Do you want to move in with me?", we can also take it as merely leading to a question like "What are your thoughts on moving in with me?" With regards to the latter, Peggy's answer in (10) is perfectly on topic; her financial situation is relevant towards the general topic of them moving in together. If Carina herself is unsure what exactly the topic was at the time of Peggy's answer, she might just take the latter to indicate that she was mistaken about it and accept it.

Moreover, *indirect* strategic subject-change is often difficult to detect. By presupposing rather than outright expressing a change of subject, hearers might focus on the contribution itself while paying less attention to what gets accommodated by it. Therefore, participants can accept contributions leading to subject-change without being aware that they are doing so.

4.2.2.3 The subject-change is not objected to

Even if their interlocutor registers a speaker's attempt to strategically change the subject, they might not object to it, and in that case, the conversational context will change in the way the speaker wants it to. Generally, not accepting a conversational contribution or blocking a presupposition from being accommodated is perceived as rude (Langton 2018; Peet 2017). Doing so indicates that the contribution was somehow "off", inappropriate, or otherwise uncooperative. So, to avoid being hostile, interlocutors might choose to go with an attempt to strategically change the subject, rather than risk being confrontational by objecting.

Such a reaction becomes even more likely in cases where the subject-change is difficult to detect. To return to *Moving in*, it is perfectly possible that Peggy merely arrived at a different interpretation of the QUD than Carina did, and that her answer is not an attempt to dodge the question, but just to respond to her understanding of it. Therefore, even if Carina thinks that Peggy is trying to distract from her question, she might not feel sufficiently confident in her belief to confront her about it.

4.2.2.4 The subject-change is rejected

Finally, even if an attempt to strategically change the subject *is* objected to, it might still be successful. To see how, it will be helpful to use *Vaccines and Parties*. In this case, Boris Johnson tries to avoid answering a question about whether he held an illegal party at his flat. By listing his alleged achievements at the time, he tries to replace the QUD “Did Johnson have a party at his flat?” with a different one, like “What was Johnson doing in Downing Street to help the country?”

Johnson’s attempt to dodge the question was not accepted. The next speaker after him accused him of “not being able to tell [his audience] where he was”, later, he was accused of “not wanting to tell the truth”. Johnson again responded by claiming that he did truthfully tell where he was, and that he answered the question.

Of course, there is much to criticize about Johnson’s answer. But purely from the perspective of subject-change, he was successful. Presumably, Johnson wanted to change the QUD such that he neither had to admit nor deny having hosted a party at Downing Street. Had he admitted it, he would have faced the fiercest possible criticism. Had he denied it, he would have been liable to consequences for lying to parliament later. But after his non-answer, the QUD did shift – towards whether he answered the question he was asked – and he avoided both unfavorable outcomes. The very fact that people objected to his attempted subject-change led to a subject-change of its own, and one that allowed Johnson to avoid the exact topic he disliked.

What this shows is that even if attempts to strategically change the subject are rejected by the audience, speakers might still be successful in bringing about subject-change. Of course, this is not always a valid strategy, as it will leave the audience frustrated and perhaps even resentful towards the speaker. But at least sometimes, it gets them where they want to be.

To summarize then, strategic subject-change can be successful in three ways: By being accepted outright, by being accepted because the audience decides not to object, or by being rejected but thereby leading to a further argument that leads away from the initial topic.

4.3 Strategic subject-change and plausible deniability

So far, I have attempted to explain what strategic subject-change is, and how it can be a successful strategy to avoid unwanted topics. While I believe this to be independently interesting, it can also help us better understand cases currently discussed in the philosophical literature. I will turn to one of these cases now, that of plausible deniability.

4.3.1 Plausible and not-so-plausible deniability

Plausible deniability, roughly speaking, is the ability of speakers to deny having meant something their audience took them to communicate. Take the following example from Lee & Pinker (2010, 785) which is commonly cited as a paradigm case (repeated here from chapter 2.3):

*Speeding Driver*¹⁸: A driver stopped for speeding says to the police officer:

(3a) I'm in a bit of a hurry. Is there any way we can settle this right now?"

When accused of offering a bribe to the police officer, the driver responds with:

(3b) I didn't mean to offer a bribe! I was just wondering if I could pay my fine right away.

The police officer takes (3a) to be an attempt to bribe her. When being called out on this, the driver responds by claiming that she did not mean to bribe the officer. Instead, in uttering (3b), she proposes a different interpretation of (3a), which is more benign.

In recent debates in the political philosophy of language, plausible deniability has been taken to explain why speakers can make objectionable utterances while avoiding blame for them. For example, Khoo (2017, 2018) and Saul (2018, 2019) argue that by communicating racist beliefs using coded language, speakers can plausibly deny having meant something racist by claiming that all they meant was the literal meaning of their utterances. Similarly, Saul (2012, 2024) and Camp (2018) argue that by misleading rather than lying, speakers retain deniability about having deceived their audiences, by resorting to the literal meaning of their utterances. Finally, as Boogart, Jansen & van Leeuwen (2020) point out, speakers will often have deniability about the literal content of their utterances by claiming that they meant something altogether different. For example, they can insist that what they said was merely ironic, a joke, or metaphorical.

At the same time, philosophers have started investigating what it is that *gives* speakers plausible deniability (Peet 2015, forthcoming; Camp 2018, Dinges & Zakkou 2023). However, coming up with a decent account has been hindered by what is sometimes called *implausible* deniability (Berstler 2019). Consider Trump's now infamous comments regarding Mexican immigrants:

"When Mexico sends its people, they're not sending their best. [...] They're sending people that have lots of problems, and they're bringing those problems with us. They're bringing drugs. They're bringing crime. They're rapists. And some, I assume, are good people."¹⁹

¹⁸ I am here using Dinges & Zakkou's (2023, 374) reconstruction of this case, since it is a little neater than Lee & Pinker's version.

¹⁹ <https://time.com/4473972/donald-trump-mexico-meeting-insult/> (Last accessed: 14.05.2024).

This statement has rightly been called out as incredibly racist. And yet, his later VP Mike Pence defended Trump's utterances like this:

There are criminal aliens in this country, [...] who have come into this country illegally who are perpetrating violence and taking American lives. He also said many of them are good people. You [Michael Kaine, in a debate against whom Pence uttered this] keep leaving that out of your quote.²⁰

This defense checks the boxes for a speaker using plausible deniability. Pence claims that rather than meaning that Mexican immigrants are generally vile, with just some exceptions, Trump merely pointed out that there was a problem with crime committed by some of them; at worst, he might have been a bit hyperbolic. Yet calling such *implausible* cases *plausible* deniability seems to be quite a stretch.

This has created a puzzle for philosophers trying to capture what plausible deniability consists of. On the one hand, leaving implausible cases out of one's account (as Dinges & Zakkou (2023) and Peet (forthcoming) do) means ignoring the kinds of cases in which deniability is most frequently discussed in the literature on political speech. On the other hand, including them (as Camp (2018) does) leads to a too permissive notion of plausible deniability that fails to differentiate between actually plausible and implausible cases. Both seem undesirable.

4.3.2 How strategic-subject change can help

In the previous section, I explained how in many cases, the speaker appears to be making use of plausible deniability while in fact intuitively not having any. In this section, I will show that we can account for such cases by thinking of them as a speaker engaging in strategic subject-change, under the guise of deniability.

To see how, let us turn to a case that is frequently cited as a case of a politician making use of plausible deniability to avoid blame:

Inner cities: In a 2014 interview, US politician Paul Ryan made the following statement:

(13a) There is a culture of laziness in our inner cities.

“Inner cities” is a codeword frequently used to covertly refer to poor African American neighborhoods in inner city areas. Understood this way, Ryan communicated that Black Americans were specifically lazy, a kind of sentiment that was unsurprisingly met with outrage. When Ryan was then confronted by a constituent of his, Alfonso Gardner, about his allegedly

²⁰ <https://www.washingtonpost.com/news/the-fix/wp/2016/10/05/pence-yes-trump-called-mexicans-rapists-and-criminals-but-you-keep-forgetting-about-the-other-part/> (Last accessed: 14.05.2024).

racist comment, Ryan defended himself, claiming that he did not mean to communicate anything regarding race.

(13b) Race has nothing to do with this. This is about getting people to work, and this is about making sure that we can actually fight poverty more effectively.²¹

Khoo (2017; 2021) argues that by using the deniability granted by the code word “inner cities”, Ryan could claim that all he meant by (13a) was what he literally said, without the racist connotation. Thereby, he dodged criticism for his statement.

But this explanation is unsatisfactory for two reasons: First, Ryan was successful regardless of his denial *not* going through. And, in fact, had it gone through, he would *not* have been successful in what he wanted. I will explain these two reasons in turn.

As Saul (2024) convincingly argues, politicians making racist comments like (13a) and receiving criticism for it face a dilemma: If they accept the criticism and apologize, this might antagonize the part of their voter base that harbors genuinely racist beliefs. If they stand by their comments, however, this might lead to resentment from the part of their voter base that does not condone racism.

Now if Ryan intended to deny that he made a racist comment, he failed. After having uttered (13b), Alfonso Gardner pushed back, accusing Ryan of making racist:

(14) “The bottom line is this: your statement wasn't true, that's a code word for ‘black’.”

(15) “If you didn't mean this, you wouldn't have said it, OK? People don't say something that they don't mean.”

But if we accept Saul's story about Ryan neither wanting to apologize, nor to stand by his utterance, he was successful. After (13b), the topic of the conversation changed from whether (5a) *was* racist to whether Ryan *meant* something racist by saying it. This allowed him to avoid the initial QUD – whether he stood by the racist implications of (5a) – altogether. Ryan was successful despite his denial not going through.

Furthermore, I think we are licensed to make an even stronger claim: If his denial *had* been successful, he would *not* have gotten what he wanted. Assume that in fact, Ryan had deniability about communicating something racist, and that he was successful in establishing that that was not what he meant. In this case, he still would have had to take a stance on whether what he said

²¹ See: <https://www.nbcnews.com/politics/first-read/paul-ryan-confronted-over-inner-cities-remark-tense-exchange-n57636> (last accessed 19.03.2024)

was okay. If the racist connotations of his original utterance were a mere accident, he should have apologized for them. No doubt his statement caused hurt and offense in some people hearing it, and if that is not what he aimed for, an apology would have been in place. Conversely, if he hadn't apologized afterwards, he would have communicated that the negative consequences of his statement were fine with him – which is akin to standing with his statement. So, if his denial had gone through (and was actually plausible), he would have ended up in the kind of situation he wanted to avoid.

We can see something similar in the Trump case discussed above. If Trump really had had deniability regarding his statement, he would have come out as *not* racist against Mexicans. But this probably would not have squared well with the racist part of his voter base. Instead, what Pence accomplished on his behalf was to change the subject – away from the claim that Mexico is *sending* immigrants, towards the claim that many immigrants engage in criminal activity. The latter is something his party is happy to endorse more generally.

We have now seen two things regarding cases of “implausible” deniability. First, if we just look at them from the perspective of deniability, they often fail. Second, by making use of implausible deniability, speakers manage to change the subject. And in doing so, they are more successful in reaching their conversational goals.

This gives us an attractive picture: Speakers who “implausibly” deny need not actually intend for their denials to go through. They may instead want to avoid the topic of their blameworthiness and thus strategically change the subject, turning the immediate QUD away from whether their utterances were morally permissible to what they meant by them. This captures why uses of implausible deniability are often successful – they are for the same reason that strategic subject-change is – and why, to quote Elizabeth Camp's (2018, 49) apt description, denial is often “so annoying”. Interlocutors are often aware that what they witness is not actual plausible deniability. But if they do not turn the conversation completely hostile, they cannot dismiss the supposed denial outright.

What this means for accounts of plausible deniability is that they can safely exclude obviously implausible cases. What explains the success of these cases is strategic subject-change, not anything related to the idea of deniability.

Still, there is something left to explain: It is often argued that rhetorical devices, such as misleading or codewords, *do* give speakers deniability. Yet, on the view that I have proposed, this is irrelevant for what is going on, as here, it is merely there changing the subject that saves them

from having to take blame. So, what do we make of the idea that these devices *do* grant deniability?

I will not be able to offer a definitive answer here, but I will make some suggestions. First, we should question what motivates the idea that speakers get deniability by using dogwhistles or the like in the first place. As far as I am aware, the main motivation for holding that dogwhistles etc. give speakers deniability is that speakers often deny having meant what they were taken to communicate with these devices (Khoo 2017; Saul 2024). But if we can explain their denial as an attempt to strategically change the subject, there does not seem to be a reason to posit that the relevant rhetoric devices give speakers deniability.

Furthermore, even if we accept that *most* speakers using codewords, misleading, insinuating and so forth do not have plausible deniability about what they meant, we do not need to hold that this is *always* the case. New codewords are created all the time (Saul 2024), and perhaps when using new ones, speakers can plausibly deny having meant the norm-violating parts of their utterance. But the more well-known a codeword is, the less plausible the denial becomes.

This brings me to my suggestion concerning whether dogwhistles and the likes grant speakers deniability: We should more carefully separate expressions that *might* grant a speaker plausible deniability from those that *actually* do. There is some plausibility to the claim that codewords *can* give a speaker deniability but holding that it always does so is quite implausible. Consider the case of the code “88” as part of someone’s license plate. “88” stands for “HH” (H being the eighth letter of the alphabet), which in turn is short for “Heil Hitler”. Now can a person plausibly deny having meant to express sympathies to Nazis when having “88” on their license plate? Intuitively, it would seem to depend. A well-known Neo-Nazi will not be able to do so, as it is very implausible that they are not aware of the coded meaning of “88”. But a person from a part of the world where Neo-nazism is a lot less salient, and who has never previously seemed to hold any right-wing sentiments? That denial would seem quite plausible.

So, it seems that we can simultaneously hold that some speakers using a codeword who deny that they meant the coded meaning do not have plausible deniability and are merely attempting to change the subject, and that codewords are the kind of rhetorical device that might grant a speaker plausible deniability, at least in certain circumstances. I think that this is a tempting idea, but I will leave a full defense of it for a different occasion.

4.4 Strategic subject-change in political speech

In this final section, I want to turn to an especially prominent area in which strategic subject-change is used: Political speech. I will attempt to give a picture of why it is in this sphere that strategic subject-change is used so ubiquitously. Given the scope of the topic, I do not aim to give a complete story about what is going on. Rather, I want to make some suggestions that warrant further, more in-depth research.

4.4.1 Politicians changing the subject: Some cases

Politicians engage in strategic subject-change all the time. As Bull (2008, 337) tentatively puts it, they are often seen as “the sort of people who will never give a straight answer to a straight question”. And it is easy enough to find ample examples of such behavior. An especially blunt example can be found in the case I have referred to as *Vaccines and parties* (see section 4.1).

Another, slightly more subtle attempt was made by Austrian chancellor Karl Nehammer, following the publication of a leaked video in which he claimed that no Austrian parents were too poor to feed their children, as they could just get them food from McDonalds. This was seen by many as not taking seriously the struggles of poor people in Austria. Nehammer defended himself, by asking

(16) Why isn't the share of people working part-time increasing - not even among women who have no caregiving responsibilities? If I don't have enough money, I go to work, but that doesn't happen.²² [Translation mine]

Rather than taking a stance towards the allegations made against him, Nehammer implied that poverty cannot be that big of an issue if more people aren't taking on extra jobs.

The cases discussed so far, including *Inner Cities* and Mike Pence defending Trump's comments about Mexico “sending their worst”, all involve subject-change as a response to a critical question. But strategic subject-change can also happen on its own, without being prompted by a demand to respond to a question.

Consider former British Home Secretary Suella Braverman's explanation of her plan to reform police in the UK. After herself bringing up the topic of police impartiality, she said that

(17) People want officers patrolling the streets, not policing pronouns.²³

²² <https://www.sn.at/politik/innenpolitik/aufregung-nehammer-video-mcdonalds-burger-mahlzeit-146039113> (last accessed 15.05.2024)

²³ *Daily Express*, October 12, 2023.

Braverman steered away from the topic from actual concerns about police participating in practices such as racial profiling towards (probably hugely exaggerated) demands for them to mind individuals' choices regarding their pronouns.

4.4.2 Why do politicians strategically change the subject?

The frequency of strategic subject-change in politics raises two questions: Why do politicians want to avoid critical questions and difficult topics? And why do they choose strategic subject-change rather than a different approach to bring that about? I will address both in turn.

In communication studies and communication psychology, it is widely accepted that politicians evade critical questions or topics to *save face* (Bull 2008). This idea goes back to the work of Erving Goffman, who characterized a person's (public) face as their reputation, authority and accepted moral standing (Goffman 1968). When speaking, individuals usually try to *maintain face* by acting in accordance with their self-perception of their own status. If they fail to live up to this, they might *lose face*; for example, if a politician demonstrates a lack of expertise on an issue, their reputation will suffer.

Empirical results suggest that politicians engage in evasive rhetorical maneuvers when they are in what is called an *avoidance-avoidance conflict* (Bavelas et al. 1990; Bull 2003). Such conflicts are characterized as situations in which “all of the possible replies to a question have potentially negative consequences, but nevertheless a reply is still expected” (Bull 2003, 130). Here, it seems that either answer will lead to a loss of face; politicians are facing a dilemma.

The cases discussed in this paper can all be characterized as avoidance-avoidance conflicts. I have already discussed above that Paul Ryan, in *Inner Cities*, most likely was unhappy with both apologizing for his comments and with standing by them, as both would have damaged his reputation in parts of his electorate. But contrary to what Bull (2003, 2008) and others write, such conflicts do not just arise in response to critical questions. Take the subject-change made by Suella Braverman in (17): Braverman was not *asked* about police impartiality, yet she brought the topic up only to change it herself. But she, too, faced an avoidance-avoidance conflict. If she had raised the issue and taken the concerns seriously, this would have antagonized people sympathetic to current UK policing practices. Meanwhile, not talking about police impartiality at all would not have gone down well with those of her potential voters concerned with such issues. Bringing the topic up but then changing the subject towards culture war issues was a middle ground between two unfavorable outcomes.

So, we have some reason to think that the primary reason for strategic subject-change in the political realm is avoidance-avoidance conflicts, and politicians wanting to save face. But this still leaves open the question of why they would pick subject-change over other evasive strategies, such as lying or misleading, not answering, or making ambiguous statements.

Let us discuss these options in turn, beginning with lying. The main reason speakers change the subject rather than lie is that in most of the cases discussed above, lying isn't really a solution to the conflict. Boris Johnson's subject-change in *Vaccines and parties* was necessary precisely because both affirming and denying that he would step down if convicted were not viable options for him. Whether either of these answers would have been truthful or a lie does not really matter for him. Subject-change helps because it completely avoids answering the question, rather than committing someone to an option, regardless of how good/bad that option is for them.

But even in cases where lying is a viable option - as one might think about the Suella Braverman comments - it still comes with risks that changing the subject does not have. If a politician is caught lying, that almost always comes with a loss of face. They might be seen as lacking integrity or expertise. Meanwhile, strategic subject-change is much harder to discover than lying, and it has a smaller stigma attached to it.

Misleading, usually taken to be indirectly communicating something deceptive (Berstler 2019, see also chapter 2), faces the very same problems as lying. Often it is not an option, as any answer to a critical question will come with a loss of face. And while misleading is generally seen as less bad than lying (which is not necessarily justified, see Saul 2012, Viebahn 2021), its discovery still comes with a loss of face that is bigger than that of strategically changing the subject.

A different strategy discussed in the literature on politicians evading questions is giving ambiguous answers. Ambiguous answers are answers that, while seemingly answering the question under discussion, do not commit the speaker to a clear evaluation of any of the answers. The classical example of this is a politician being asked about prior failures and answering with "I would do a lot differently now". While the speaker acknowledges that parts of their past actions were wrong, they do not say anything about what exactly was, or why.

Ambiguous answers fair better with regards to avoidance-avoidance conflicts, as they do not commit the speaker to any of the unfavorable options. But their ambiguous nature still leads to a less than ideal outcome. First, hearers with clear opinions on what the speaker should say in response to a question might still see ambiguous answers as missing the mark. They want a specific answer, and they do not get it. Moreover, ambiguous answers come with the risk of being disambiguated unfavorably: Hearers might interpret "I would do a lot differently now" as

actually admitting to specific wrongdoing. In this case, the speaker got a result they wanted to avoid.

Finally, simply not answering the question again comes with severe loss of face. Here, the evasive behavior is obvious, and the speaker is again judged as lacking integrity or expertise. We can see this by the negative attention Boris Johnson received at a different instance when asked about the party he hosted during a lockdown. His complete refusal to say anything with regards to the criticism made against him led to mockery, new criticism, and an *Independent* article with the very unflattering headline “Boris Johnson refuses 17 times to answer questions about Downing Street lockdown parties”²⁴. Hardly an outcome that can be described as “saving face”.

To summarize then, politicians engage in strategic subject change when confronted with an avoidance-avoidance problem which threatens their “face”. And strategic subject-change has advantages compared to other strategies one could go for in such a situation.

4.4.3 Why strategic subject-change is successful in political contexts

In the last section, I have discussed reasons for why politicians might choose strategic subject-change over other evasive rhetorical strategies. But to fully understand why subject-change is so ubiquitous in public discourse, there is another factor we need to consider: That the contexts in which politicians speak are often such that strategic subject-change succeeds very easily. In this section, I will give some reasons for why that is.

Above, I have argued that there are three ways in which strategic subject-change can succeed: the audience accepts the subject-change, the audience does not accept the subject-change but goes along with it nevertheless, and the subject-change leads to a further disagreement, which itself is a departure from the original topic. As we will see, the contextual factors discussed below can make all these outcomes more likely.

First, the conversations politicians are involved in are frequently *time pressed*. Interviews, Q&A’s, or town hall meetings follow a tight schedule. This increases the likelihood of the interlocutors just going with the subject-change. Repeating their questions or challenging the speaker’s subject-change would cost them valuable time, which they might rather spend on different questions. But the time-constrained nature of political conversations also has an impact on cases in which the interlocutors do push back on a speaker changing the subject. Even if they were able to win the ensuing argument, this just might not happen, as their time for speaking is up.

²⁴ <https://www.independent.co.uk/news/uk/politics/boris-johnson-lockdown-party-interview-sue-gray-b2019116.html> (last accessed 16.05.2024)

Second, there is often a significant difference in authority between politicians and the people they are talking to. For many circumstances, politicians have formal authority; they are the ones dictating, say, what gets talked about in a rally or a Q&A. But very often, their authority is informal: They are important people, and challenging their speech is akin to disrespecting, or at the very least questioning that authority.

In cases involving formal authority, the same consequences apply as in the time-constrained scenario. As they can dictate how much time is spent on a question, they can force others to go along with their subject change, or silence those unwilling to do so.

With informal authority, things are a little different. First, informal authority can lead audiences to accept strategic subject-change they would otherwise find questionable. Just like a star-struck fan might accept everything their idol says, a die-hard political supporter will happily perform the type of mental gymnastics required to accept the subject-change. Second, since objecting to subject-change is akin to questioning their authority, audience's might be more willing to go along with a change of subject than they would otherwise be. Authority aggravates the unwillingness to be confrontative which I have discussed above.

Finally, political conversations often happen asynchronously. What this means is that rather than being constituted by the kind of immediate back and forth typical of everyday conversations, they often happen in temporally extended and disunified ways. One politician makes a statement during an interview. Another calls in a press conference to criticize that statement. The first politician addresses the criticism when confronted with it in an interview two days later... If during such an exchange, one of the politicians involved performs strategic subject-change, calling them out on it will be very difficult; and thus, participants will be more likely to go along even if they would not want to do so otherwise.

Together with the results from the previous sections, this gives us a compelling explanation for why politicians so frequently change the subject: They have good reason to do so. First, they often find themselves in situation in which any topical response might lead to damaging their reputation. Second, in such situations, subject-change as a strategy is preferable to alternatives. And third, when they try to change the subject, they are often in a uniquely good position to do so.

4.4 Conclusion: Strategic subject change and deception

In this paper, I have tried to provide a picture of what strategic subject-change is. I have discussed some of its central characteristics, embedded it within broader theories concerning

communication, and explained how it can succeed. Furthermore, I have shown that understanding it can help us better approach other issues in the social/political philosophy of language. Finally, I have tackled a domain particularly prone to strategic subject-change – political speech – and tried to explain why it is so ubiquitous here by drawing from the previous insights.

In doing so, I have discussed all sorts of cases in which speakers might want to change the subject – cases in which they are uncomfortable talking about a subject, in which they try to avoid taking blame, or in which they prefer changing the subject. In this final section, I want to connect my findings to the overarching topic of this thesis, deception.

I will discuss three potential upshots this chapter has for broader philosophical debates about deception. The first concerns the connection between deception and plausible deniability. As I have discussed in chapter 2.3, plausible deniability is frequently taken to be able to explain how speakers can deceive their audience, and then avoid the blame that usually comes with deception: By plausibly denying that they intended to communicate the deceptive content (see Saul 2012b, Khoo 2021).

If what I have argued for in this chapter is correct, this picture, at least in many cases, does not hold. What often looks like a speaker using plausible (or not-so-plausible) deniability is in fact them strategically changing the subject – away from the moral status of their speech, towards what exactly it was they said (or meant to say). So, at the very least, we need a broader picture of how speakers avoid the culpability that comes with deceiving, one that includes changing the subject.

A second upshot that is closely related to the first concerns the motivation speakers have with regards to misleading rather than lying. Some philosophers and linguists have argued that one of the prime motivations for this is that misleading grants a speaker deniability where lying does not (see Camp 2018, Pinker & Lee 2010, Saul 2012b; see also chapter 2.3). But again, if their supposed deniability is (at least often) better explained by them changing the subject, this cannot be the full story. Something else is needed to explain why speakers mislead rather than lie, at least in many cases.

Finally, recognizing that strategically changing the subject is often an alternative to deceiving ones audience (as can be seen in *Moving In*, but also in cases of politicians using strategic subject-change) raises important new questions regarding the moral status of deception. While philosophical debates have often compared lying to misleading in this regard (see chapter 2.2), we should also ask ourselves if changing the subject is in fact always morally preferable to such means of deception, as is apparently frequently assumed. Doing so might very well lead to a more nuanced view of the morality of deception.

5. Conclusion

In this thesis, I have tried to make two main contributions to the current philosophical landscape concerning deception. First, I have proposed a novel argument regarding the moral status of the lying-misleading distinction, arguing that misleading is often morally worse than lying. Second, I have argued that speakers frequently try to avoid culpability for morally problematic speech-acts, such as lying or intentionally misleading one's audience, by strategically changing the subject. Furthermore, I have shown that recognizing this calls into question the relevance of plausible deniability for avoiding culpability, which goes against what is currently assumed in the literature.

There is a lot I couldn't touch on in this thesis. For example, I have been largely silent about the issues discussed in chapter 2.1, that is, the rhetorical devices used for lying and misleading, respectively. For the purposes of chapter 3, I have adopted minimal conditions for lying and misleading that are largely in line with the received view on both. It would be interesting to investigate whether adopting a commitment-based view like Viebahn's (2021) would lead to a different outcome regarding the moral status of the lying-misleading-distinction.

Likewise, my brief exploration of strategic subject-change has left many questions open. For example, I have not proposed any strict definition of strategic subject-change; nor have discussed whether having one would be helpful in the first place. Furthermore, I have not touched on the moral status of strategic subject-change. Intuitively, many of its uses are morally objectionable. Boris Johnson dodging questions regarding his potential illegal partying habits certainly is not the correct way to go about things. But on closer inspection, things are not quite that simple. We might for example think that in many cases, strategically changing the subject is preferable to telling someone an uncomfortable truth. Thus, moral considerations regarding subject-change certainly merit a closer investigation.

Still, I hope that I have managed to raise some interesting new questions and problems regarding deception. If nothing else, writing this thesis was quite interesting in its own right, and also a lot of fun.

6. Literature

- Adler, J. E. (1997): "Lying, Deceiving or Falsely Implicating", *The Journal of Philosophy* 94/9, 435-452.
- Almagro, M. & Torices, J. (2018): "The Nature of (Covert) Dogwhistles", in Saborido, Oms & González de Prado (eds.): *Proceedings of the IX Conference of the Spanish Society of Logic, Methodology and Philosophy of Science*, 93-100.
- Augustine [395] (1952): „Lying“, in *Treaties on Various Subjects* Vol. 16, edited by R. Deferrari. New York: Fathers of the Church, 53-120.
- Bull, P. (2003): *The Microanalysis of Political Communication*. Claptrap and Ambiguity. London: Routledge.
- Bull, P. (2008): "‘Slipperines, Evasion, and Ambiguity’: Equivocation and Facework in Noncommittal Political Discourse", *Journal of Language and Social Psychology* Vol. 27 Iss. 4, 333-344.
- Baumann, H. (2015): "Gibt es einen moralisch relevanten Unterschied zwischen Lügen und Irreführen?", *Zeitschrift für Praktische Philosophie* 2(1), 9-36.
- Bavelas, J.B., Black, A., Chovil, N., Mullett, J. (1990): *Equivocal communication*. Newbury Park: Sage.
- Bennett, C. (2018): "The Alteration Thesis: Forgiveness as a Normative Power", *Philosophy & Public Affairs* 46 No. 2, 207-233.
- Berstler, S. (2019): "What's the Good of Language? On the Moral Distinction between Lying and Misleading", *Ethics* 130, 5-13.
- Boogart, R., Jansen, H., Van Leeuwen, M. (2020): "‘Those are Your Words, not Mine!’ Defence Strategies of Denying Speaker Commitment", *Argumentation* 35, 209-235.
- Camp, E. (2018): „Insinuation, Common Ground and the Conversational Record“, in: Fogal, Harris, Moss (eds.): *New Works on Speech Acts*. Oxford: Oxford University Press, 40-67.
- Carson, T. (2006): "The definition of lying", *Noûs* 40, 284-306.
- Chisholm, R., Feehan, T. (1977): "The Intent to Deceive", *The Journal of Philosophy* 74(3), 143-159.
- Coleman, L., Kay, P. (1981): "Prototype Semantics: The English verb ‘lie’", *Language* 57, 26-44.
- Davies, A. (2016): "How to Silence Content with Porn, Context and Loaded Questions", *European Journal of Philosophy* 24(2), 498-522.

- Descartes, R. [1644] (2017): *Principles of Philosophy*. Translated by John Cottingham. Edited by Jonathan Bennett. <https://www.earlymoderntexts.com/assets/pdfs/descartes1644part1.pdf>
- Dinges, A. & Zakkou, J. (2023): “On Deniability”, *Mind* 132, 372-401.
- Fallis, D. (2009): “What is Lying?”, *The Journal of Philosophy* 106(1), 29-56.
- Frankfurt, H. (2002): “Reply to G. A. Cohen”, in *Contours of Agency: Essays on Themes from Harry Frankfurt*. Cambridge: MIT Press, 340–344.
- Fricker, E. (2012): “Stating and Insinuating”, *Proceedings of the Aristotelian Society, Supplementary Volumes* 86, 61-94.
- Goffman, E. (1967): *Interaction ritual: Essays on face to face behavior*. Garden City: Anchor.
- Grice, H. P. (1975): “Logic and conversation”, in *Studies in the Way of Words*. Cambridge (MA): Harvard University Press, 41-58.
- Holton, R. (2019): “Lying About”, *The Journal of Philosophy* 116, 99-105.
- Isenberg, A. (1973): “Deontology and the Ethics of Lying”, in *Aesthetics and Theory of Criticism: Selected Essays of Arnold Isenberg*. Chicago: University of Chicago Press, 245–264.
- Kant, I. [1797] (1949): “On a Supposed Right to Lie from Altruistic Motives”, in *Critique of Practical Reason and Other Writings in Moral Philosophy*, translated by L.W. Beck. Chicago: Chicago University Press, 346-350.
- Kant, I. (1963): *Lectures on Ethics*, translated by L. Infield. New York: Harper.
- Keller, P. (forthcoming): “Testimony by Presupposition”, *Erkenntnis*.
- Khoo, J. (2017): “Code Words in Political Discourse”, *Philosophical Topics* 45 (2), 33-64.
- Khoo, J. (2021): “Code Words”, in: Khoo, Sterken (eds.): *The Routledge Handbook of Social and Political Philosophy of Language*. New York: Routledge, 147-160.
- Lackey, J. (2013): “Lies and deception: an unhappy divorce”, *Analysis* 73, 236-248.
- Langton, R. (2018): “Blocking as Counter-Speech”, Fogal, Harris, Moss (eds.): *New Work on Speech Acts*, Oxford: Oxford University Press, 144-164.
- Lee, J. & Pinker, S. (2010): “Rationales for Indirect Speech: The Theory of the Strategic Speaker”, in *Psychological Review* 117/3, 785-807.

- Lefevere, J., Goovaerts, I., Turkenburg, E. (2023): “Changing the Subject: An Analysis of Candidates’ Issue Emphasis in Televised Election Debates, 1985-2019”, *Journal of Broadcasting & Electronic Media* 67(5), 755–774.
- Lewis, D. (1975): “Languages and Language” in Gunderson (ed.): *Language, Mind, and Knowledge*. Minneapolis: University of Minnesota Press, 3-35.
- Lewis, D. (1979): “Scorekeeping in a Language Game”, *Journal of Philosophical Logic* 8(1), 339-359.
- Lo Guercio, N. & Caso, R. (2022): “An account of overt intentional dogwhistling”, *Synthese* 200 (3), 1-32.
- MacIntyre, A. (1994): “Truthfulness, Lies and Moral Philosophers: What Can We Learn from Mill and Kant?”, *Tanner Lectures on Human Values*, delivered at Princeton University.
- Marsili, N., Löhr, G. (2022): “Saying, Commitment and the Lying-Misleading-Distinction”, *The Journal of Philosophy* 119(12), 687-698.
- Moore, J. G. (2000): “Did Clinton lie?”, *Analysis* 60, 250-254.
- Peet, A. (2015): “Testimony, Pragmatics, and Plausible Deniability”, *Epistémè* 12/1, 29-51.
- Peet, A. (2017): “Defective Contexts”, in: Khoo, Sterken (eds.): *The Routledge Handbook of Social and Political Philosophy of Language*. New York: Routledge, 193-207.
- Peet, A. (forthcoming): “The Puzzle of Plausible Deniability”, *Synthese*.
- Pepp, J. (2022): “What is the Commitment in Lying?”, *The Journal of Philosophy* 119(12), 673-686.
- Pinker, S., Nowak, M., Lee, J. (2008): “The Logic of Indirect Speech”, *Proceedings of the National Academy of Sciences* 105/3, 833-838.
- Roberts, C. (2012): “Information structure in discourse: Towards an integrated formal theory of pragmatics”, *Semantics & Pragmatics* 5, 1-69.
- Rees, C. F. (2014): “Better lie!”, *Analysis* 74, 59-64.
- Saul, J. (2012a): “Just go ahead and lie!”, *Analysis* 72, 3-9.
- Saul, J. (2012b): *Lying, Misleading, and What is Said: An Exploration in the Philosophy of Language and in Ethics*. Oxford: Oxford University Press.
- Saul, J. (2018): “Dogwhistles, Political Manipulation, and Philosophy of Language”, in: Fogal, Harris, Moss (eds.): *New Works on Speech Acts*, Oxford: Oxford University Press, 360-383.

- Saul, J. (2019): “What is happening to our norms against racist speech?”, *Aristotelian Society Supplementary* Vol. 93, 1-23.
- Saul, J. (2024): *Dogwhistles and Figleaves*. How Manipulative Language Spreads Racism and Falsehood. Oxford: Oxford University Press.
- Simpson, D. (1992): “Lying, Liars and Language”, *Philosophy and Phenomenological Research* 52, 623-639.
- Sorensen, R. (2007): „Bald-faced lies! Lying without the intent to deceive”, *Pacific Philosophical Quarterly* 88, 251-264.
- Sliwa, P. (2020): “Excuse without Exculpation. The Case of Moral Ignorance”, in Shafer-Landau (ed.): *Oxford Studies in Metaethics Volume 15*. Oxford: Oxford University Press, 72-95.
- Stalnaker, R. (1999): “Assertion”, in: *Context and Content*. Oxford: Oxford University Press, 78-95.
- Stalnaker, R. (2002): “Common Ground”, *Linguistics and Philosophy* 25, 701-721.
- Strudler, A. (2010): “The Distinctive Wrong in Lying”, *Ethical theory and moral practice* 13(2), 171-179.
- Stokke, A. (2013a): “Lying and Asserting”, *The Journal of Philosophy* 110(1), 33-60.
- Stokke, A. (2013b): “Lying, Deceiving, and Misleading”, *Philosophy Compass* Vol. 8 Iss. 4, 348-359.
- Stokke, A. (2016): “Lying and Misleading in Discourse”, *Philosophical Review* 125/1, 83-134.
- Stokke, A. (2018): *Lying and Insincerity*. Oxford: Oxford University Press.
- Timmermann, F., Viebahn, E. (2021): “To lie or to mislead?”, *Philosophical Studies* 178 (5), 1481-1501
- Viebahn, E. (2020): “Lying with Presuppositions”, *Nous* 54/3, 731-751.
- Viebahn, E. (2021): “The Lying-Misleading Distinction: A Commitment-Based Approach”, *The Journal of Philosophy* 118/6, 289-319.
- Viebahn, E. (2022): “Lying, Misleading, and Fairness”, *Ethics* 132, 736–751.
- Viebahn, E. (forthcoming): “What does it take to tell a lie?”, in Wiegmann (ed.): *Lying, Fake News, and Bullshit*. London: Bloomsbury, 1-24.
- Walton, D. (1996): “Plausible Deniability and Evasion of Burden of Proof”, *Argumentation* 10, 47-58.

Webber, J. (2013): “Liar!”, *Analysis* 73(4), 651-659.

Williams, B. (2002): *Truth and Truthfulness: An Essay in Genealogy*. Princeton: Princeton University Press.

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