

A Misleading Tendency

by

Nadia Miller

A thesis

presented to the University Of Waterloo

in fulfilment of the

thesis requirement for the degree of

Master of Arts

in

Philosophy

Waterloo, Ontario, Canada, 2023

© Nadia Miller 2023

Author's Declaration

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

Communication is a powerful tool. We use it to share information, express how we feel, coordinate in cooperative endeavors, and so much more. However, we regularly use standard communicative devices like speech and gesture to piss each other off, to manipulate each other, and to cause harm; flipping you off or calling you by a slur is no less clear a way of sending you a message than is stating a simple fact. In roughly the last decade, philosophers of language have turned their attention away from idealized communicative contexts to develop new theories about non-ideal communication, which includes phenomena like lying, misleading, and deceptive speech. The philosophical study of lies and deception predates the discipline of philosophy itself, but contemporary theories place a special emphasis on their explanatory and moral applications in real-world social and political contexts. While care has been taken in this tradition to define lying speech, and to distinguish it from speech that is ‘merely’ misleading or deceptive, insufficient care has been taken to define and distinguish misleading speech as a distinctive communicative act. This gap in the literature overlooks a class of cases that are not captured by definitions of lying and deception and their unique moral significance.

In this thesis, I defend the following, novel account of misleading:

Misleading: A communicative act C is misleading iff C has a tendency to cause others to reason badly

I argue that this definition of misleading is the best way to capture the fact that misleading does not necessarily cause anyone to form a false belief. My view challenges prior accounts of misleading in the literature which take this element to be necessary, and thereby fail to accommodate cases of misleading that are not successful or result in something other than a false belief.

Acknowledgements

This research was funded by the Social Science and Humanities Research Council of Canada (SSHRC)'s Canada Graduate Scholarships – Master's program.

I am especially grateful to my supervisor, Professor Jenny Saul, for her support and feedback during this project and throughout my graduate career. I am also thankful to my parents, Fiala and Sebastian, and to my siblings, Yousef and Lara. My best ideas are born over family dinners.

Table of Contents

AUTHOR’S DECLARATION	II
ABSTRACT	III
ACKNOWLEDGEMENTS	IV
INTRODUCTION	1
MEANING AND METHOD	2
<i>Concepts</i>	2
<i>The Conceptual Project</i>	4
CHAPTER 1: THE LITERATURE ON LYING, DECEPTION, AND MISLEADING	7
LYING	7
<i>Non-Warranting Context</i>	8
<i>Lying in misbelief</i>	9
<i>Irony, hyperbole, and metaphor</i>	10
<i>In defense of intent</i>	11
DECEPTION.....	15
<i>The intent to deceive is not necessary</i>	16
<i>On “causing”</i>	18
MISLEADING– ?	19
<i>Merely Misleading</i>	20
<i>Misleading</i>	23
CONCLUSION.....	26
CHAPTER 2: THE MISLEADING-DECEPTION DISTINCTION	28
IS MISLEADING THE SAME AS DECEIVING?.....	28
<i>The Law on Misleading and Deception</i>	33
MISLEADING AND DECEPTION: A MORAL DIFFERENCE?	37
CONCLUSION.....	41
CHAPTER 3: A MISLEADING TENDENCY	42
MISLEADING: MY DEFINITION	42
<i>Misleading without causing false beliefs</i>	44
<i>Misleading without success</i>	46
<i>Objection: Isn’t everything misleading?</i>	47
WHAT’S WRONG WITH MISLEADING?	50
LYING, DECEPTION, AND A MISLEADING TENDENCY	54
CONCLUSION	56
REFERENCES	57

Introduction

All those minds that are interested in finding out the truth communicate with each other across the distances of space and time. I, too, was taking part in the effort which humanity makes to know.

– Simone de Beauvoir, *Memoirs of a Dutiful Daughter* (1959)

Communication is a powerful tool. We use it to share information, express how we feel, coordinate in cooperative endeavors, and so much more. Beauvoir describes what many take to be the ideal function of communication, one that is self-evidently valuable. However, just as evident is the fact that communication oft goes awry from this ideal. We regularly use standard communicative devices like speech and gesture to piss each other off, to manipulate each other, and to cause harm; flipping you off or calling you by a slur is no less clear a way of sending you a message than is stating a simple fact. In roughly the last decade, philosophers of language have turned their attention away from idealized communicative contexts to develop new theories about non-ideal communication, which includes phenomena like lying, misleading, and deceptive speech. The philosophical study of lies and deception predates the discipline of philosophy itself, but today's contemporary theories place a special emphasis on their explanatory and moral applications in real-world social and political contexts. While care has been taken in this tradition to define lying speech, and to distinguish it from speech that is 'merely' misleading (Saul, 2012; Stokke, 2018) or deceptive (Carson, 2010), insufficient care has been taken to define and distinguish misleading speech as a distinctive communicative act. This

gap in the literature overlooks a class of cases that are not captured by definitions of lying and deception and, in doing so, overlooks their unique moral significance.

In this thesis, I defend the following, novel account of misleading:

Misleading: A communicative act C is misleading iff C has a tendency to cause others to reason badly

I argue that this definition of misleading is the best way to capture the fact that misleading does not necessarily cause anyone to form a false belief. My view challenges prior accounts of misleading in the literature which take this element to be necessary, and thereby fail to accommodate cases of misleading that are not successful or result in something other than a false belief.

In Chapter 1, I discuss the background philosophical literature on lying, deception, and misleading. By the end of the chapter, I consider two positions on misleading: that misleading is either the same as deception, or that they are different in that deception is immoral whereas misleading is not necessarily. In Chapter 2, I argue that neither of these alternatives suffices to capture what misleading is. In Chapter 3, I propose and defend my novel account of misleading.

Meaning and Method

Concepts

In the analytic spirit, my strategy will be to investigate and hopefully to explain what conditions are necessary and sufficient for something to count as an instance of misleading. The goal of this strategy in general, which is called ‘conceptual analysis,’ is to clarify a concept that is otherwise unclear or imprecise. The term ‘concept’ is itself in need of clarification, but that is a project far beyond the scope of mine here. For our purposes, a ‘concept’ is something like a Fregean *sense* or Carnapian *intension* (with an ‘s’). On a dyadic theory of meaning originating in

Gottlob Frege and continuing through the work of Rudolf Carnap and other logical empiricists, the semantics of an expression are divided into its extension and its intension. Roughly, the extension is the set of things that a term refers to; for example, the extension of “Ottawa” is Ottawa, and the extension of “red” is all the objects to which the term can truly applied; i.e., all things that possess the property of *redness*. Technically, an intension is a function from possible worlds to extensions. For our purposes, it is fine to think of the intension of a term as something like a detailed definition which, if understood by a speaker, could be used as a fairly reliable rule for deciding whether an object satisfies the concept. Terms with different intensions may be co-extensive; that is, they can differ in meaning even if they refer to the same thing. Take the example of Yousef, who is both a sports coach and a chef. The players he coaches know him only as ‘the coach’ and the cooks in his kitchen know him only as ‘the chef.’ Though all the players and the cooks are referring to the same entity—Yousef—when they use ‘Coach’ or ‘Chef,’ the terms nevertheless differ in meaning: for one thing, it makes perfect sense for Joan, a cook, to say both that she knows the chef and that she does not know the coach; for another, if Joan decided to join Yousef’s team, it seems she has genuinely discovered a new fact, that *the chef is the coach*. In this case, despite referring to the same thing, it is clear that ‘coach’ and ‘chef’ differ in what they mean.

An intension may just be a singular concept, represented by a singular term, or it may be a proposition made up of constituent parts including concepts and other grammatical items, which together can be represented as a sentence. A simple way of explaining a concept’s intension is to provide a list of its necessary and sufficient conditions. For example, to be a human, or to apply the concept correctly to some entity *X*, it is sufficient that *X* is a rational being, (but not necessary, given that we count babies and people with cognitive disabilities as

humans). On the other hand, it is necessary that X is a mammal for X to be human— there are no humans that are non-mammal, and a person who insisted otherwise would appear to be misusing the term. (Though, it is not sufficient; many mammals are not humans.) Necessary and sufficient conditions let us do three things: They let us make truth evaluations, distinguish felicitous from infelicitous uses of language, and individuate metaphysical objects. As an example of the first, consider a condition stating that, for a thing to be water, it is both necessary and sufficient that it has a molecular makeup of two hydrogen atoms to one oxygen atom. Thus, the truth of the statement “ Y is water” can be evaluated empirically by figuring out whether Y is made up of H_2O molecules, (or analytically if Y is synonym of “water”). Felicitous use of the term “water” is constrained by the same conditions; so, a person who says “ Z is water,” where Z is made up of C_2H_6O , appears to have made one of two mistakes, with the latter of which being linguistic in nature: Either they are mistaken about the entity they are describing— they think they have a Y and not a Z before them— or they are mistaken about the meaning of “water,” and the fact that it is only to be applied to Y s. Finally, with this necessary and sufficient condition for water, we have a rule for making ontological distinctions between water and other chemical objects; we can label objects in our vicinity by following intensional rules.

The Conceptual Project

Traditionally, conceptual analysis was used by philosophers of science who were concerned with the concepts belonging to mathematical and scientific languages, such as Carnap and other members of the 20th Century Logical Empiricist tradition. One of the core tenets of the Logical Empiricist movement was a commitment to using clear and precise language; namely, language that is made up for the most part by truth-evaluable constructions which could be arranged grammatically to form premises and arguments. In other words, propositions.

Ordinary language functions quite differently to mathematical and scientific language. For one thing, communication, the cornerstone of our intelligent existence, is ripe with vagueness, ambiguity, irony, metaphor, error, misunderstanding, gestures, insinuations and a million other things that make it far from an ‘ideal’ subject of analysis. Many of the empirical cases discussed here arise in more or less ordinary contexts, but my target concept is not supposed to replace ordinary understanding. Instead, my suggestion is that philosophers, and specifically philosophers of language, with an interest in misleading should prefer to use the definition of misleading that I propose here.

My decision to use cases from the real world and imaginary cases that illicit ordinary linguistic intuitions is purposeful. One of Carnap’s desiderata for conceptual explication is *similarity*, which states that ‘explicandum,’ or the concept after it has been explained, must capture more or less of what it was originally supposed to capture in its extension. If the conclusions of a conceptual analysis are drastically different (for example, by entailing that many or most speakers are mistaken in their understanding of the concept), then these differences should raise at least some suspicion. This is not to say that every single ordinary usage of the term ought to be accounted for by the explicandum— it may be that ordinary speakers are sometimes mistaken, and there is no doubt that much philosophical terminology is specialized— often to the verge of idiosyncrasy— but unless there is reason to think that people who use the term are completely missing the mark, the extensions of the original concept and the explicandum (or what I will also call the ‘target concept’) should overlap. Although misleading does not necessarily involve speech, linguistic data will be the primary evidence base for my research. The phenomenon I am writing about is communicative in nature, thus my definition will refer to misleading ‘communicative acts’ rather than misleading acts in general. The

information communicated may be about the world, or it may be about the mental states of the person attempting to communicate something. In the cases I will look at, even when speech is not involved (such as when people mislead through omission or silent action), there still remains the basic structure of at least two people inferring meaning from each others' behaviour.

The goal of my argument is to persuade a reader with a philosophical interest in misleading to accept two things: First, that there is a problem or confusion with how philosophers are thinking about misleading; and second, that my proposal is the best conceptual characterization of misleading. The evidence for my argument is made up of a variety of ordinary cases in real-world contexts like political speeches and courtrooms; hypothetical scenarios designed to illicit intuitions, and ideas from the contemporary (*2000 onwards*) philosophical literature on lying, misleading and deception. If my conclusions have any prescriptive force, it is intended towards philosophers of language and others who theorize about misleading. It is not assumed to extend to the concept of misleading ordinarily used and understood by most English language users— indeed, I am not convinced that I am even in a position where I might influence this group in the first place. My target audience, so to speak, are my academic peers.

Chapter 1: The Literature on Lying, Deception, and Misleading

The concepts ‘lying’, ‘misleading’ and ‘deception’ denote phenomena with at least two things in common: A tendency to arise in linguistic discourse, and a tendency to prevent people from acquiring true beliefs. This chapter is the beginning of an attempt to pick apart these three concepts, and to explain how they relate to each other; to communication; and to our moral intuitions about them. The chapter will culminate with a philosophical puzzle: Is misleading different to deception, and if so, how?

Lying

My mother, Fiala, hides the baking chocolate in a not-so-secret kitchen cabinet. We are never allowed anywhere near the baking chocolate but, yesterday, I ate it all. I was in the kitchen when she came down to bake chocolate chip cookies this evening, aware that it was only a matter of time before she asked, “Nadia, did you eat the baking chocolate?” At this moment, I knew that she knew that any of her three children could have eaten the chocolate, so I uttered (1) to deceive her into believing that it was not me:

(1) No, I did not eat the baking chocolate

Most would agree that I told a lie, but what is it about the scenario that makes this so? One line of reasoning goes as follows: I said something false; therefore, I told a lie. However, if saying something false is sufficient for lying, then the following scenario would also amount to lying: 5 minutes earlier, Fiala asked my sister, Lara, if she’d help bake the cookies. Before agreeing, Lara asked:

“Do we have baking chocolate?” To which, Fiala responded:

(2) There is baking chocolate in my secret cabinet

As a matter of fact, (2) is false— there wasn't chocolate in the cabinet, because I ate it. Therefore, if saying something false is sufficient for lying, then Fiala lied. However, I am confident that the reader will share my intuition that Fiala did not lie by uttering (2). This intuition can be explained by the fact that Fiala did not know that her utterance was false. If this is right, then we can amend our previous characterization of my lie in (1): I said something *which I knew to be false*, therefore, I lied. This is susceptible to at least three sorts of counterexample: Saying something known to be false in a non-warranting context; saying something believed to be false which happens to be true; and irony, hyperbole, and metaphor.

Non-Warranting Context

Morgan Freeman knows that he is not Nelson Mandela; likewise, Daniel Craig knows that he is not James Bond. After his performance in *Invictus*, it would be strange to accuse Morgan Freeman of lying when he introduced himself as Nelson Mandela onscreen, or Daniel Craig for saying “The name’s Bond, James Bond” instead of “The name’s Craig, Daniel Craig.” In a context where speakers are acting or, say, performing a stand-up set (where it is perfectly acceptable for— and expected that— comics will invent and embellish stories), the audience is generally aware that the truth of what is being said is not guaranteed.

The concept of warrant is central to Carson’s definition of lying in *Lying and Deception: Theory and Practice* (2010). According to Carson, one warrants the truth of their statement when one “promises or guarantees, either explicitly or implicitly, that what one says is true” (26). On his account, one lies just in case they say something that they know to be false, in a context where they warrant the truth of their statement.¹ It is hard to say when, where, and how often it is

¹ A person S tells a lie to another person S1 iff: 1. S makes a false statement X to S1, 2. S believes that X is false or probably false (or, alternatively, S does not believe that X is true), 3. S states X in a context in which S thereby warrants the truth of X to S1, and 4. S does not take herself to be not warranting the truth of what she says (29–30)

that one actually finds oneself in such a context. A paradigm example would be a courtroom where one is a witness under oath, but we are rarely under oath in ordinary life. I agree with Carson's observations that it would be pointless to try and explicate necessary and sufficient conditions for what counts as a warranting context, and that whether one is in a warranting context will depend largely on conventional and cultural norms about when people are expected to be truthful, and when people know that they are expected to be truthful. In any case, it seems that one will struggle to lie if one's audience already expects them to be untruthful.

Lying in misbelief

According to Carson, a person must not only believe that they are saying something false, but the statement must be false in actuality. Saul (2012) questions whether saying something false is necessary for lying (p. 5–6). In her example, Tony believes there are no weapons of mass destruction in Iraq, but would like to convince people that there are. He utters (3):

(3) There are weapons of mass destruction in Iraq

Even if it turns out that there are weapons of mass destruction in Iraq, the fact that Tony had no idea makes (3) a lie. I am inclined to agree, though I recognize that this brings about conflicting intuitions. One might argue that leaving out the requirement that a lying statement be false undermines the moral significance of lying; in other words, if lying is wrong because it brings about false beliefs in others, then a lie which brings about a true belief either is either not a lie, or is not the kind of lie that we should worry about. However, this is just one way of thinking about what is wrong lying. As an alternative, Shiffrin (2021) contends that lies are wrong because they abuse the channel of testimony as a “reliable and distinctive mode of access to the contents of the speaker's mind” (p. 269). If we agree with Shiffrin, then the moral significance of lie should not

turn on whether it misrepresents the world, but rather on whether it misrepresents the speaker's mental states.

Irony, hyperbole, and metaphor

The final counterexample to lying as saying something known to be false is neatly illustrated by examples of irony, hyperbole and metaphor. Imagine Fiala finds out that I ate the chocolate and, rolling her eyes, utters (4) and then (5):

(4) How *uncharacteristic* of you

(5) That was a ton of chocolate

In fact, (4) is false— it is totally characteristic of me to eat any and all chocolate in the vicinity, and my mother knows this best. (5) is also false: there were two Hershey's packets at about 250g each and, being a careful baker, Fiala knew their precise measurements. Yet, under the circumstances, it would be strange to characterize either of these utterances as lies; Fiala was clearly using irony in her utterance of (4), and hyperbole in her utterance of (5).

It may be tempting to say that, using irony and metaphor, Fiala thereby caused the context to become a non-warranting one. This is wrong, for if she was not warranting the truth of her statements, then it would be very difficult for us to make sense of them as ironic/metaphorical. In a counterfactual scenario where both Fiala and I know that it was not characteristic of me to eat the baking chocolate, (4) would not pass as irony; and had it only been a measly amount of chocolate, (5) would be false on both its literal and metaphorical interpretation. For ironic and metaphorical utterances to succeed as such, they must be truth-sensitive in a way that fictional utterances need not be. I take it that this is partly explained by the fact that irony and metaphor are nevertheless used to communicate information about the actual world, whereas fiction is designed to do quite the opposite.

Before saying a thing or two about communication, it would be helpful to consolidate the preceding observations with our working definition of lying. For now, I propose to adopt a variation of Saul's definition of lying in her 2012 *Lying, Misleading and What is Said*:

If the speaker is not the victim of linguistic error/malapropism or using metaphor, hyperbole, or irony, then they lie iff (1) they say that P; (2) they believe P to be false; (3) they take themselves to be in a warranting context. (p. 19)

In defense of intent

Standard definitions of lying require that the speaker possess a deceptive intent. This condition usually amounts to something like a *mens rea* for lying; A cannot lie to B unless A intends to deceive B. This is compatible with a definition of lying that does not require A to make a false statement—just as I can be guilty for attempted murder even if no one died, I can be guilty for attempting to lie even if no one believes me. The morally relevant factor is that I *intended* for someone to believe me. However, Saul and Carson reason that examples of bald-faced lying warrant rejection of this principle. Consider a version of Carson's example:

A man on the witness stand in a courtroom has witnessed a murder. Because there is CCTV footage that clearly shows the man witnessing the murder, and this footage has been presented to the jury, everyone knows that everyone knows that the man saw the crime take place. But, for fear of reprisals, when asked whether he saw the murder, the witness says,

I did not see the murder.

The intuition that the witness is lying is clear, and yet it is equally clear that he is not intending to deceive anyone.²

The witness knows that no one will believe their lie, so it is hard to see how they could intend to deceive— at most, they could intend to present themselves as having not seen the murder, but this would be quite different. The intent to deceive implies that the deceiver at least thinks it is possible to make others believe them, but this element is missing from the case. Nevertheless, we should like to say that the witness told a lie.

Aside from *lying under duress*, we may also think about cases of *forced lies*. To use Saul's example, a malicious state leader may force his people to chant false political slogans which they do not believe in by threatening to kill or imprison those who refuse (p. 10).

Though I agree that the intent to deceive in itself is not necessary for lying, there are reasons to preserve at least some degree of intent as a necessary condition. Saul's definition does this by precluding linguistic error and malapropism from counting as lies. For example, imagine I meant to say that I danced the Flamenco on Friday but accidentally said,

(6) I danced the flamingo on Friday

Most would be hard pressed to say that (6) was a lie, even though I did not believe that I took a flamingo dancing. I accept Saul's qualifications, but I think further qualifications are in order to exclude other non-lying cases such as the following: A person with Alzheimer's says something that they otherwise believe is false, and takes themselves to be warranting the truth of their utterance. Or, a person suffering from a delusional episode utters something that they know is false. I contend that what is said must proceed from the agency of the speaker to count as a lie;

² Carson (2010, 20)'s example is amended by Stokke (2013, 349)

which is not to say that they must intend to say something that they believe is false, only that they intend to say what they said. This requirement may be phrased as ‘knowingly’ saying that P.

If we accept this agential requirement for lying, then we bring about the happy implication that cases of linguistic error and, more generally, non-deliberate speech will not count as lies. In the case of malapropism, I did not know that I was saying something about flamingos– if I did, then I could rightfully be accused of lying. The ‘knowingly saying’ requirement also takes care of Saul’s humorous case of accidental falsehood:

A nice example comes from a colleague's time on a post-doc in Mexico. Anna, an English rock climber, wanted to tell her colleagues that many people in England climb without ropes. So she uttered [U]:

[U] En Inglaterra hay mucha gente que escala sin ropa.

[U] actually means that in England there are many people who climb without clothes. This claim is false, but Anna did not lie; she accidentally said something false, through a linguistic error.

(p. 14)

Thus, I propose the following amendment to Saul’s definition:

Lying: If the speaker is not using metaphor, hyperbole, or irony, then they lie iff [1] they *knowingly* say that P; [2] they believe P to be false; [3] they take themselves to be in a warranting context.

Going forward, when I talk about lies and lying, I am talking about lying. Let us briefly review the listed conditions. A person lies iff they

[1] Knowingly say that P

[2] They believe P to be false

[3] They take themselves to be in a warranting context

Special attention should be paid to the notions of ‘saying’ invoked in this definition. To say that P, a person must produce a sentence whose literal meaning expresses the proposition P. The idea of ‘literal meaning’ is hotly contested in philosophy of language. On a minimalistic conception of literal meaning, what is said is taken to be determined by the conventional meaning of words and their syntactic combination, permitting a certain and limited degree of contextual enrichment as is required to disambiguate, or to fix the referent of indexical and demonstrative terms like “I,” “this,” “here,” and so on. The biggest problem for the minimalist view as far as lying is concerned is that lying requires the person to express a proposition, P, but the P that liars express is often underdetermined by the words they actually uttered. Returning to our first example, when my mother asked if I had eaten the baking chocolate, I replied with (1):

(1) No, I did not eat the baking chocolate

(1) expresses a complete proposition, namely, ‘I did not eat the baking chocolate.’ The proposition is expressed ‘completely’ because it admits of a truth evaluation. However, I may have uttered (1*) in response to my mother, which does not express a complete proposition:

(1*) No

Taken in isolation, the word “no” cannot be evaluated as true or false. And yet, we should still like to say that I lied if I uttered (1*). That (1*) nevertheless expresses the same proposition as (1) is explained by the context of the utterance; that is, I uttered (1*) in response to a question and, drawing on this extra-linguistic information, we may ‘complete’ my utterance to be a literal expression of (1). Saul captures the need for completion using a ‘Needed for Truth-Evaluability’ criterion (NTE):

(NTE)A putative contextual contribution to what is said is a part of what is said only if without this contextually supplied material, S would not have a truth-evaluable semantic content in C.

(p. 61)

It is difficult to specify the extent to which we really draw and rely on contextual resources to determine what someone has said. In contrast to the minimalist and NTE view, some have argued that contextual enrichment determines literal meaning far beyond disambiguation, referent-fixing, and propositional completion. On the most radical contextualist account, there is no difference between what is said and whatever meaning is interpreted by hearers. Withholding judgement on which view is correct in general, for our purposes here, it will not do. I will return to the issue of literal versus non-literal communication in the final section on misleading but, for now, let us turn to deception.

Deception

Some but not all lying amounts to deception, and not all deception counts as lying. Lying amounts to deception just in case the liar intends to cause someone to form a false belief and succeeds in doing so. These concepts come apart in four places: First, as we have discussed already, not all lying involves a deceptive intent; second, not all lies are successful. Carson's example of the witness's bald-faced lie illustrates both: The witness neither intends to make people believe that they did not see the murder, nor do they cause anyone to believe that they did not see the murder. Third, lying is about a relationship between what is said and the speaker's mental states— hence, it is not necessary for a liar to say what is *in fact* false, only that they believe that they are saying something false. This distinction may be summarized succinctly:

Lying is characterized by a relationship between what is said and a speaker's mental states, whereas deception is characterized by a relationship between one person and another's ability to get at the truth *itself*. Finally, unlike lying, one need not *say* something false to deceive. Indeed, deception may be altogether non-linguistic; for instance, a non-Muslim woman could wear a hijab with the intention of making others believe that she is a Muslim woman; if others form this belief, then she has deceived them. Though I am sensitive to non-linguistic cases, this is a project in the philosophy of language, so these will not be central to my focus.

The intent to deceive is not necessary

The preceding account of deception accords more or less with ordinary usage of the term: A deceives B by doing X iff [intent] they intend to cause B to form a false belief by doing X, and [success] B actually forms a false belief as a result of A's doing X.³ That intent plus success are sufficient to deceive is uncontroversial, but 'deception' is sometimes understood just as 'causing a false belief.' On its face, this is too broad; if I sigh calmly while slipping into an ice-cold pool because I find it soothing, and someone infers from my sigh that the pool is warm in temperature, it would be strange to say that I thereby deceived that person. Thus, both the deceptive intent and the success condition may appear to be necessary. Though I agree that unsuccessful cases are at most *attempts* to deceive, and thus that success must be necessary, I believe there are moral reasons to reject deceptive intent as a necessary condition.

There are examples in U.S law where deception need not be intentful. Shiffrin (2019) describes two compelling cases. The first is an advertisement for Anacin-3 which stated that 'Hospitals recommend acetaminophen, the aspirin-free pain reliever in Anacin-3, more than any other pain reliever.' The statement was literally true, acetaminophen was the pain-killer most

³ Versions of this definition are endorsed by Carson (2010) and Fallis (2018).

commonly prescribed by hospitals, but as Shiffrin describes, “the advertisement had the tendency to lead audience members to draw the false conclusion that Anacin-3 was the brand prescribed most by hospitals.” (p. 73). In a second example from the California Court of Appeals,

A car dealer was live-broadcasting a commercial on his car lot, as he made a regular practice of doing. Defendant, consumed with fury because the dealer had parked a car on his adjoining property, stormed up to him while he was on air, in full view of the cameras, and menacingly hurled an epithet at him implying canine parentage. Although Defendant was just an angry neighbour, the plaintiff dealer contended that Defendant’s speech and behaviour conveyed to viewers and potential customers that Plaintiff had cheated Defendant in a prior car sale. The jury found Defendant liable for slander; an appellate court agreed that a reasonable audience would incorrectly take Defendant to be an aggrieved customer protesting Plaintiff’s unfair business practices. Indeed, a member of the television audience testified that from watching Defendant’s outburst, she ‘came to the conclusion that if one bought a car from respondent, one might expect unfavorable treatment’. (p. 75)⁴

What makes these cases of deception despite the absence of deceptive intent, in contrast with the non-deceptive case of my sighing in the pool? They are deceptive, Shiffrin would say, because the defendant neglected to consider that their behaviour was susceptible to misinterpretation. After all, it is reasonable to assume that, in some contexts, people have a duty to ensure that their statements are not misinterpreted, especially where harm could foreseeably ensue. Commercial statements, such as those made in advertisements or at investors meetings, are cases where this duty legally obtains. However, it seems that a weaker version of this duty is also present in many ordinary contexts. Consider: I am approached by a lost tourist in my city,

⁴ *White v. Valenta*, 234 Cal. App. 2d 243, 245 (1965)

who asks me for directions to the nearest Starbucks. There are two big white buildings at left turns ahead, one leads to Starbucks and the other leads to a long, winding path into the country. I say “left at the big white building, then your first right.” Here, the foreseeable risk that the tourist will become even more lost if they turn at the wrong building plausibly suggests that I have at least some moral duty to make sure that the tourist knows exactly what I mean by my statement. If I foresaw this risk, and my utterance caused them to form the false belief that they should turn at the wrong building, then it seems fair to say that I negligently deceived them. On the other hand, such a duty does not seem to apply to my sighing in the pool; in this case, the person’s forming a false belief was accidental to any predictable consequence of my behaviour.

On “causing”

So much for deceptive intent, but what does it mean to “cause” a false belief, and is this really necessary for a behaviour to count as deceptive? Shiffrin writes that deception always involves the acquisition or reinforcement of a false belief as a non-accidental response to the deceiver’s behaviour (p. 71). I do not suppose that her use of ‘non-accidental’ implies intent, but rather that it implies direct relationship between the deceptive behaviour and the epistemic or deliberative processes of the deceived. Other terms used to describe the causal relationship between the deceptive behaviour and subsequent false belief include ‘drives’ and ‘causes to persist’. This raises questions about deceptive reinforcement and deceptive omissions. Imagine, after discovering the missing baking chocolate, my mother forms the belief that Lara had in fact eaten it, and makes this belief transparent to me by muttering “Lara can’t be trusted around the chocolate” under her breath. I may reinforce her false inference by uttering (7):

(7) *I know right!* She really can’t be trusted around the chocolate.

Though I cannot be said to have caused Fiala's false belief by uttering (7), I deceived her insofar as I have caused her pre-existing false belief to persist. Furthermore, had I stayed silent instead of uttering (7), I could foresee that my silence would cause my mother's false belief to persist.⁵ Thus, I would have deceived her by omission. Given these considerations, I propose to accept Shiffrin (2019)'s view of deception as a working definition:

Deception: forming a false belief (or having one reinforced) due, in substantial part, to the relevant agency of another person—whether through her qualifying speech, silence, behaviour, or an omission. (p. 71)

Misleading— ?

To the best of my knowledge, the most substantial pieces of philosophical literature that treat misleading as a distinctive phenomenon are Saul (2012) and Stokke (2018). Still, Saul and Stokke only treat one kind of misleading in depth, which Saul refers to as *merely* misleading. On Saul's theory, one merely misleads if they communicate that which they believe to be false without literally saying it. To distinguish merely misleading from lying, Saul primarily uses examples of conversational implicature, where a speaker implicates a proposition they believe to be false while saying something that they believe to be true. Her hallmark example is of Bill Clinton during an interview where he denied having a sexual relationship with Monika Lewinsky by uttering, in the present tense,

(8) There is no improper relationship

⁵ For the most part, Fiala has the reasonable expectation that her children will own up to their mistakes

Clinton knew that, if (8) was uttered at a time t , and if there was no improper relationship between them at t , then (8) was literally true. Thus, he would not have lied. However, given that the topic of the interview was about his past relationship with Lewinsky, Clinton may be reasonably interpreted as implicating that there was *never* an improper relationship between him and Lewinsky. Indeed, the nature of the topic suggests that this would be *the most* reasonable interpretation of (8), since he was expected (and knew he was expected) to discuss the history of their relationship as opposed to its status in that very moment.

Merely Misleading

Saul's conception of Merely Misleading employs the technical notion of what is said discussed earlier, which draws on no more contextual information than is needed to determine a truth evaluable proposition (NTE). The sense in which speakers may mislead by merely 'implicating' that which they believe to be false is also a technical notion in the philosophy of language. In his 1975 *Logic and Conversation*, Grice coins the notion of a 'conversational implicature.' In simple terms, a conversational implicature is a proposition that a listener must retrieve in order to interpret the speaker as making a co-operative contribution to the conversation. Grice captures this with his co-operative principle

(CP): "Make your conversational contribution such as is required, at the stage at which it occurs, by the accepted purpose or direction of the talk exchange in which you are engaged," (p.45)

which he breaks down into four categories:

Quantity: Make your contribution as informative as is required (for the current purposes of the exchange); Do not make your contribution more informative than required (p. 45)

Quality: Try to make your contribution one that is true

Relevance: Be relevant

Manner: Be perspicuous (p. 46)

Grice's maxims are not to be thought of as a normative guide to conversation; rather, they describe principles that people tend to follow when they are trying to maximize cooperative endeavour in a communicative exchange. Occasionally, we bend these rules on purpose and with the intention of being recognized as doing so. Metaphor involves a deliberate flouting of the maxim of quality, and to be interpreted correctly, we need our audience to realize that we mean something quite different to what we have said. Take, for example, when Fiala said (5):

(5) That was a ton of chocolate

(5) is obviously false— I take for granted that few people keep a ton of chocolate in their homes, yet alone that a single person could eat such an amount (though I am open to the challenge).

Thus, (5) is a clear violation of Grice's maxim of quality. To preserve the thought that my mother is partaking rationally in the conversation, I must supply an alternative interpretation which makes sense of her utterance as a co-operative contribution. Thus, I can infer that she really means (5*):

(5*) That was *a lot* of chocolate

Misleading via implicature involves saying something believed to be true, while violating Gricean maxims in such a way as to implicate something believed to be false. In the context of Clinton's interview, (7) was at least a violation of Relevance, insofar as the interview was clearly about a past relationship while the utterance was about the existence of a relationship *right now*. It was probably also in violation of Quantity for the same reason, and Manner for how obscure it sounded. Thus, to preserve the charitable supposition that Clinton was trying to be maximally co-

operative during the interview, a reasonable audience member could easily interpret him to mean (7*):

(7*) There was never an improper relationship

Another clearcut example of misleading by implicature arose during the Prime Minister's Questions (PMQs) session in the UK House of Commons on the 18th December 2020. During the session, then PM Boris Johnson was being questioned about the COVID regulation violations that took place during the 'Partygate' scandal at 10 Downing. Johnson stated:

(9) I repeat that I have been repeatedly assured since these allegations emerged that there was no party and that no covid Rules were broken. That is what I have been repeatedly assured.⁶

In the final HoC report on the matter, the committee found:

By claiming that when he referred to having been repeatedly assured, by "repeatedly" he had meant merely "on more than one occasion". We note that this is contrary to common English usage. It is clear that when Mr Johnson used the term "repeatedly" at PMQs, he wished his audience to suppose that there had been multiple occasions at which assurances had been given, rather than merely more than one, and, as suggested by our evidence, possibly as few as two. (p. 56)

On a strictly literal interpretation, Johnson's statements do indeed mean "on more than one occasion," which would make them true in the case that he had only been assured twice. As it

⁶ Matter referred on 21 April 2022 (conduct of Rt Hon Boris Johnson): Final Report; Committee of Privileges, UK House of Commons.

turned out, there was conclusive evidence of at most two such assurances that were made to Johnson, and sworn testimony on behalf of two of his peers that there were more (only they could not say whether they were 100% sure, and could recall neither who made the purported assurances, nor when). However, were he following Grice's maxim of quantity, he should have said "I was twice assured." By saying "repeatedly" (and by saying it repeatedly), the most reasonable interpretation of his statement would be that he was assured many times over, and certainly more than just twice. Unlike lying, merely misleading has the advantage of preserving the speaker's deniability; it allows one to convey something false while precluding lying accusations. If they are called out for implying something that they believed to be false, the misleader may simply invoke the literal meaning of their utterance and rebut, *I didn't say that!* And this is precisely what Johnson did in his own submission to the committee report:

This statement was entirely accurate, and I do not believe that the House has been misled by it.

By "repeatedly" I meant on more than one occasion and by more than one person. (p. 40)

So, to reiterate, a person merely misleads when they communicate something they believe to be false, without literally saying it.

Misleading

Merely misleading is a purposely constrained notion of misleading. It's primary function is to clarify the muddy distinction between literally saying something believed to be false, and communicating something false via pragmatic content. These two phenomena are closely related and hard to distinguish in ordinary discourse, since both tend to be undertaken with similar aims (such as to deceive), and have similar outcomes (a misrepresentation of the speaker's beliefs). However, understanding misleading in this way is limited in that it only characterizes cases of

misleading that are structurally similar to lying in at least two respects: First, the cases of merely misleading we have discussed were all deliberate attempts to convey something that the speaker believed was false; that is, they are all cases of *knowingly* misleading. Second, they all involved the communication of a proposition. However, misleading need not be knowing; as I will argue in Chapter 2, it is not necessary for a person to be aware that they are misleading for us to say that they misled. Moreover, misleading does not necessarily involve the communication of a proposition; as I will argue in Chapter 3, a person need only cause someone to reason badly in order to mislead, and it is possible to do this without causing them to stand in any relation to a particular proposition.

Insufficient care has been taken by philosophers of language to characterize misleading as such. There is some discussion of misleading in legal theory, and lots of technical discussion in courtroom dossiers (especially with respect to commercial matters like advertising and investor securities). These sources will prove useful in figuring out what misleading really is— or at least what is supposed to be— however, legal usage is often idiosyncratic and not bound to conform to ordinary usage of the term. So, I will leave the discussion of legal notions of misleading for the proceeding chapter. Now, let us consider what else philosophers have said about misleading.

To distinguish misleading from lying and deception, first recall the definitions from preceding sections, as well as a tentative definition of merely misleading:

Lying: If the speaker is not using metaphor, hyperbole, or irony, then they lie iff they [1] *knowingly* [2] say that P; [3] believe P to be false; and [4] take themselves to be in a warranting context.

Deception: Forming a false belief (or having one reinforced) due, in substantial part, to the relevant agency of another person—whether through her qualifying speech, silence, behaviour, or an omission. (Shiffrin, 2019, p. 71)

Merely Misleading: Communicating something believed to be false without literally saying it

In addition to merely misleading, Saul (2012) makes some observations about misleading in general:

‘Misleading’, unlike ‘lying’, is a success term. A has not misled B unless B believes A. A can, however, lie to B even if B does not believe A. In order to pinpoint the issue correctly, successful lying is what should be compared to misleading; and lying should be compared to attempting to mislead. (p. 71)

Thus, according to Saul’s observations, success is a necessary condition for misleading. However, this is the only condition she offers, and if success alone is both necessary and sufficient⁷ for misleading, then we will be hard pressed to find a distinction between misleading and deception.

Carson preserves the intent to deceive in his definition of deception for a ‘practical’ reason: the term “deception” usually carries with it a negative evaluative connotation that makes

⁷ Saul does not say whether she takes the success condition to be sufficient.

little sense if deception is not inherently intentional (p. 47). He therefore reserves “mislead” for cases where someone inadvertently causes another to form a false belief. He writes:

The negative evaluative connotations of the term "deception" are inappropriate if we hold that there can be completely unintentional deception (for which people are often blameless). There is an important distinction to be drawn between intentionally and unintentionally causing others to have false beliefs. Rather than coin new words, we can use the terms "deceive" and “mislead” (or "inadvertently mislead") to mark this distinction. **Roughly, to mislead someone is to cause her to have a false belief.** It is not self-contradictory to say that someone misled another person unintentionally or inadvertently. (p.47)

Thus, Carson defines misleading: To mislead is to cause a false belief, intentionally or unintentionally. However, only a very strong understanding of moral blameworthiness entails the conscious intent to commit an immoral action. Both legal and conventional standards for psychological blame are easily established by demonstrated knowledge that one is committing the immoral action, even if it is not one’s conscious goal. Reckless behaviour with reasonably foreseeable consequences also amounts to *mens rea* and, although it is admittedly a much weaker standard, so is negligence. A large chunk of the next chapter will be dedicated to picking apart the different elements of psychological blameworthiness.

Conclusion

So far, we have established that the characteristic feature of lying is saying something believed to be false in a context where one is taken to be, and takes oneself to be, warranting the truth. This means that lying is always about a relationship between what we believe, what we say, and where we say it. We also considered that lying sometimes but not always involves the

intent to deceive, and that even deception does not necessarily involve deceptive intent. Yet, while we have certainly made ground on the kind of misleading that is formally nearest to lying—that is, *merely* misleading—we find ourselves only a little less perplexed about what the distinction between misleading and deception could be. At this stage, we are left with two obvious choices: Either, misleading is identical to deception, or, misleading amounts to deception that is not wrongful. In the next chapter, I will argue that neither choice suffices to capture what misleading really is. Then, in the final chapter, I will attempt to find an option that does.

Chapter 2: The Misleading-Deception Distinction

In the last chapter, we defined deception as “Forming a false belief (or having one reinforced) due, in substantial part, to the relevant agency of another person—whether through her qualifying speech, silence, behaviour, or an omission.” (Shiffrin, 2019, p. 71). While the lying-misleading distinction has been thoroughly theorized, by the end of the chapter, it seemed like the distinction between misleading and deception was less clear. It turns out that philosophers define the misleading-deception distinction in one of two ways: Either, the concepts of misleading and deception are used interchangeably, implying that they are the same thing; or, they are morally different. The goal of this chapter is to see if what philosophers have to say about the misleading-deception distinction can help us get to the essence of misleading.

Is misleading the same as deceiving?

A deceives B about a proposition p iff A successfully causes B to believe p or reinforces B’s pre-existing belief in p . A set of necessary and sufficient conditions for this definition would be success, [1] B believes something false; and causality, [2] A was the cause of B’s belief or belief reinforcement. [2] appears as though it presupposes [1]—it is hard to see how one can cause/reinforce something unless they succeed at doing it. However, the success condition only says that, if A deceived B about a proposition p , then B believes p . [2] is required to establish the causal link between A’s behaviour and B’s forming that belief or having it reinforced. Moreover, without [2], the definition would risk including any person who succeeds at causing or reinforcing something *other* than a false belief in the minds of their audience. For example, A may succeed at eliciting a negative emotion in B’s mind without necessarily deceiving B about a fact.

A common practice in the philosophy of language is to use the terms “mislead” and “deceive” interchangeably. As Andreas Stokke explains in his 2018 *Lying & Insincerity*:

Writers on lying often use mislead and deceive interchangeably, and sometimes the distinction we have just noted is described as a difference between lying and other forms of linguistic deception, or the like. I follow this practice insofar as I will not be concerned with proposing distinctions in meaning between these words. (I tend to use mislead, misleading, etc. when speaking of the contrast between lying and deceptive speech that is not lying.) (p.77)

Thus, according to Stokke’s definition, a theory of merely misleading is just a way of capturing non-lying deceptive speech. Saul takes up a similar position on misleading. She writes that, for something to count as misleading proper, a success condition like [1] must be satisfied. As we saw in Ch1.3:

‘Misleading’, unlike ‘lying’, is a success term. A has not misled B unless B believes A. A can, however, lie to B even if B does not believe A. In order to pinpoint the issue correctly, successful lying is what should be compared to misleading; and lying should be compared to attempting to mislead. (p. 71)

Whether Saul takes causality to be essential to misleading is not clear, but we can infer that A’s behaviour must be at least a partial cause of B’s belief– it would be difficult to say that A misled B if B’s belief was caused by something entirely irrelevant to A’s behaviour.

I think that this is an issue about metaphysical identity as well as the interchangeability of words– which is to say, I believe that philosophers’ views on lying, misleading, and deception have implications for how we think about the phenomena themselves, not just the meaning of

terms. Though co-referential terms may differ in meaning, synonymous terms always co-refer, and I struggle to see how a person could use terms interchangeably if they did not take them to be synonymous. From this, we may assign metaphysical commitments to a philosopher even if they do not make them explicit: If they use the terms “misleading” and “deception” without a semantic difference, then we should be reasonable in our assumption that they believe misleading and deception are the same thing. Call this position *Misleading-Deception Identity* (MD-I):

MD-I: Misleading and deception are the same

At least Stokke is committed to MD-I. Someone who accepts Stokke’s point that merely misleading amounts to non-lying deceptive speech could alternatively use “misleading” to denote a *form* of deception or vice-versa. This would not have the same metaphysical implication as MD-I. One can use either or both of “primate” and “mammal” to talk about primates, but the statement “primates are mammals” is not trivial because “primate” and “mammal” are not synonymous; primates are a kind of mammal, and all that is true of mammals must be true of primates, but not the other way. Likewise, if misleading is a form of deception, either term can be used to talk about misleading but not deception. The decision to use one and not the other would not be trivial in the way that it would be if MD-I were true. It would involve a weaker, unidirectional similarity between misleading and deception. Either:

MD-S1: Misleading is a kind of deception

Or

MD-S2: Deception is a kind of misleading.

There are a few possible strategies for establishing that MD-I, MD-S1, and MD-S2 are wrong. MDI is wrong if there is disagreement between the necessary and sufficient conditions for

misleading and deception; that is, if a condition *C* is necessary or sufficient for one of misleading or deception but not both, then MD-I is false (Leibniz's law). MD-S1 is wrong if there is a *C* that is necessary or sufficient for deception and not for misleading. MD-S2 is wrong if there is a *C* that is necessary or sufficient for misleading and not for deception.

Unlike deception, misleading does not necessarily cause a *false* belief. For example, a person can merely mislead about something factually true as long as they believe it to be false. Imagine I am dogsitting and suddenly have reason to believe that the dog has escaped. To avoid lying, I utter 10 when the owner calls to check in:

(10) She's not downstairs at the moment, she could be sleeping

Knowing that I implicated (10*)

(10*) The dog is upstairs in her bed

If, unbeknownst to me, the dog turned out to be upstairs all along, or if she turned out to be safe anywhere else in the house (which would probably be another reasonable interpretation), it nevertheless stands that I merely misled by my utterance of (10). This is one reason to think that [1] is not necessary for misleading.

To be sure, let us say for now that [1] is necessary for misleading. Recall the Johnson example and imagine that, by this stage in his career, no one believes Johnson about anything anymore, including (9). Johnson's utterance does not count as a lie since he did not literally say something that he took to be false. And, since deception implies success, it would thus be wrong to say that he deceived; the most we can say is that he attempted to deceive. Finally, since we are assuming for now that success is necessary for misleading as well, then it would follow that Mr. Johnson did not mislead either.

To further illustrate, I imagine a variant of Carson (2010) and Stokke (2013)'s bald-faced lying scenario. Bald-Faced Misleading Witness:

A man on the witness stand in a courtroom has witnessed a murder. Because there is CCTV footage that clearly shows the man witnessing the murder, and this footage has been presented to the jury, everyone knows that everyone knows that the man saw the crime take place.

In Bald-Faced Misleading Witness, the witness truthfully testifies that he has a vision impairment which leaves him effectively blind unless he is wearing glasses or contacts. Moreover, in the CCTV footage, it is clear that he was wearing glasses. For fear of reprisals, when asked whether he saw the murder, the witness says,

(10) I wasn't wearing my contacts.

The witness has not perjured himself as long as (1) is true, even though it implicates a falsehood, (10*):

(10*) I couldn't see the crime

As Saul explains, it would be the lawyer's duty to extract a direct answer from the witness by asking the right questions. (Presumably, in terms of legal strategy, the misleading witness would not get very far if he avoided directly answering questions). Moreover, I do not see a straightforward way of establishing that the misleading witness attempted to deceive anyone, given that he knows that everyone else knows (10*) is false (hence, the misleading witness is *bald-faced*). The witness did not lie, nor did he successfully cause anyone to believe his implicature. If success is necessary for both misleading and deception, then the intuition that the

witness did something misleading and/or deceptive must be wrong. The conclusion that Johnson and the witness neither lied, nor misled, nor deceived is deeply counterintuitive.

If [1] is necessary for deception to take place, but not misleading, then both MD-I and MD-S1 are incorrect: Misleading is neither identical to deception, nor is it a form of deception. At most, misleading can be part of an attempt to deceive but, even if this is *often* what misleaders do, the relation is still much weaker than MD-I or MD-S1: At most, it implies that the two sometimes occur in conjunction.

The Law on Misleading and Deception

The legal distinctions between lying, misleading and deception are not clear in every context. However, the Federal Trade Commission (FTC) typically does not require proof that someone believed the defendant to establish that they misled. I do not take the law to be a decisive authority, but it serves as useful data because conceptual analysis is commonplace in the courtroom. Legal professionals are constantly required to define, redefine, and apply technical concepts, to pick them apart and interpret their ambiguities. It would be unwise to ignore decades of literature scrutinizing the very concept I am trying to define. Rules about false, misleading and fraudulent statements are central to, for example, commercial law, defamation law, and investor securities.

First, the law is not particularly concerned with lies, misleading, or deception where the message turns out to be true. This is not reason to discard the intuition that, for instance, Tony lied when he said (3) because he thought it was false, even though it happened to be true that there were weapons of mass destruction in Iraq without his knowing; or, that I lied when I merely misled by saying (10) even if my belief that the dog had escaped was false. My best guess as to why the law is not so concerned with accidental truths is that there would not seem to

be any predictable harm arising from the dissemination of truths. A corporate representative who lies about the price of shares without knowing that they had in fact come to that price in effect did nothing illegal, but still told a lie. So an important distinction to draw between legal and other philosophical discussions about lying, misleading, and deception is that the law talks about false statements rather than lies.

Commercial law about deceptive advertising does not require proof that someone was caused to believe something false in order to establish that misleading or deception has taken place. Petty and Kopp (1995) provide a six-part strategic analysis of advertising challenges, defined as “claims by consumers, competitors, or regulatory bodies [such as the FTC] that an advertisement is either false, misleading, or deceptive” (pp. 1). The six parts of any challenge are initiation, interpretation, deception, verification, remediation, and intention.

When a challenge is initiated, the first step is for the court to establish what the message of the advertising is. Unlike perjury law, which turns strictly on literal meaning, the interpretation stage of a DA claim takes a more pragmatic approach by interpreting the message “as a whole” and “in context,” to include anything reasonably implied, suggested, or otherwise conveyed. Evaluating the claim in-context also depends on who the target audience of the advertisement is supposed to be (interpreting a message intended for children, for example, would be a very different task than if it targeted adults).

From Petty and Kop’s analysis, the interpretation stage is essential and indeed nearly sufficient to establish misleading or deception:

If the advertising has been interpreted to contain a false or misleading claim, it is almost certain that deception will be found, and if the verification stage (next) confirms falsity, a

remedy will be imposed. For implied claims, deception is generally presumed; for explicit claims, the presumption is very strong. (pp. 4)

Technically speaking, proving deception should require the plaintiff to show that they believed the falsehood, that it was materially relevant to their financial decision-making, and that they relied on it to make a financial decision. However, in 1984, the Federal Trade Commission replaced the believability, materiality, and reliance elements with the requirement that the advertisement was “likely to mislead” or that it had “the capacity to mislead.” Thus, proof that consumers *in fact* formed a false belief is not required to establish that an advertisement was misleading or deceptive; “to condemn misleading advertising neither of these statutes requires proof that consumers were actually misled” (pp. 4). As long as the courts interpret a false or misleading message in the ad, the potential defense that no one believed the claim is likely to fail.⁸

This supports my point that success is not a necessary condition for misleading. However, this point was only a part of my effort to show that misleading and deception are not the same. Both deception and misleading can be established in court without proving success, so, it might be that misleading and deception are the same thing as far as the law is concerned. The reader should find this unsatisfying, for even if they are convinced by my attempt to show that misleading and deception (as defined by Shiffrin) have different necessary conditions, they may disagree with my decision to use Shiffrin’s definition. On a definition of deception where [1] and [2] are not necessary, it could still be consistent to maintain MD-I or MD-S1.

⁸ The advertiser can claim puffery as a defense (e.g., “RedBull gives you wings!”, to show that the claim wasn’t ‘likely to mislead,’ but once it has been established, the attestation that no one *in fact* believed is not a defense against deception. (Petty and Kopp, 1995, pp. 6).

I defended my decision to use Shiffrin's definition of deception, but Chapter 1 also saw Carson (2010)'s definition as an option, whereby the intent to deceive was essential. In fact, Carson distinguishes lying and misleading on exactly those grounds: misleading, unlike deception, is often inadvertent. (DA law also tends to maintain a position similar to Carson's, only without his success condition.)

Let us cover this base by postulating an alternative definition of deception and seeing how MD-I (Misleading and deception are the same) and MD-S1 (Misleading is a form of deception) hold up. Leaving aside success for now, say that the intent to deceive is both necessary and sufficient for deception. Is intent necessary for misleading? Certainly it is not. Bald-Faced Misleading Witness is one example. Consider another: Rather unfortunately, I share my birthday with the anniversary of the Capitol Hill riots. Imagine Lenny, who knows nothing about me, joins late in the conversation I am having with Benny on January 7th 2021. While raving about winning a game of mini golf on my birthday, I say (11):

(11) Yesterday was the best day ever— I stormed in and totally dominated! I am already looking forward to doing it again next year.

Lenny nervously infers from (11) that I was a participant in the riots, and confronts me later for being a Trump radical. Lenny is very relieved to find out that I am not a Trump radical, however, it would be out of place for him to complain that I had deceived him—I wasn't even talking to him in the first place, and why on earth would I wish to be known as a Trump radical? On the other hand, it would be acceptable for Lenny to say that I misled him by failing to realize how my utterance would be interpreted and not clarifying what I meant. The best way to explain this is to say that I misled Lenny by mistake, but this would be a contradiction if the intent to deceive was an essential element of misleading. So, whether one accepts the success-based or the intent-

based definition of deception, maintaining either of MD-I or MD-S1 has some weird and counterintuitive implications.

Misleading and deception: A moral difference?

So far, we have addressed whether MD-I and MD-S2 are plausible characterizations of the relationship between misleading and deception. The first, Misleading-Deception Identity, states that misleading and deception are the same thing. I inferred the metaphysical nature of this commitment via some claims about language use: namely, that a person who uses two terms synonymously is committed to the sameness of their referents. Then, I argued that both MD-I and MD-S2 are implausible because essential features of deception (like success and intent) are not essential features of misleading.

The final way that philosophers sometimes draw a distinction between misleading and deception is to say that deception is immoral whereas misleading is not necessarily. This view is expressed by Carson (2010).

The negative evaluative connotations of the term "deception" are inappropriate if we hold that there can be completely unintentional deception (for which people are often blameless). There is an important distinction to be drawn between intentionally and unintentionally causing others to have false beliefs. Rather than coin new words, we can use the terms "deceive" and "mislead" (or "inadvertently mislead") to mark this distinction. Roughly, to mislead someone is to cause her to have a false belief. It is not self-contradictory to say that someone misled another person unintentionally or inadvertently. (p.47)

First, I interpret Carson to be committed to MD-S2; namely, that deception is a form of misleading. On his view, a person misleads iff they cause a false belief; and since these two elements are always present on his account of deception, deception is just misleading with intent.

This way of defining misleading disagrees with some of my conclusions from earlier, such as that misleading need not be successful (illustrated by Bald-Faced Misleading Witness), and that it need not be the cause of a *false* belief (illustrated by examples [3] and [10]). Moreover, in the last chapter, I argued that Shiffrin's definition of deception was best equipped to handle moral and legal intuitions suggesting that the intent to deceive was not necessary, and that deception sometimes reinforces a pre-existing false belief rather than causing a new one.

Moreover, Carson's view does not explicitly account for the fact that, in the courtroom *and* in ordinary life, intent is not the sole determinant of whether an action is immoral. For example, while Shiffrin's definition does not involve the intent to deceive as such, it does require a certain degree of psychological culpability on behalf of the deceiver— something akin to *mens rea*. Analogously, the intent to kill is not necessarily present in the mind of someone who is morally blameworthy for killing; for instance, most would agree that killing someone while driving drunk is deplorable without even raising the question of whether the driver formed the conscious intention to do it. We also hold bad doctors liable when their patient dies as a result of medical negligence. Recklessness, negligence, and knowing the risk of one's behaviours are all subsumed under the Latin heading of *mens rea*, along with intent.⁹ Thus, we may interpret the moral distinction between misleading and deception in a few ways. The first is to agree with Carson, who says that deception is intended, while misleading is not necessarily. On this view, deception excludes cases where a person lacks intent but is conscious of the risk that their behaviour may reinforce or cause another's false belief, as well as cases where a false belief is caused or reinforced unconsciously— all of which would be more aptly labelled as misleading.

⁹ This is how *mens rea* is defined under the influential Model Penal Code, 1962.

Call this view MD-M1. The second is to say that deception is conscious whereas misleading is not necessarily (MD-M2). Let us lay these out and see how they fare:

MD-M1: Deception is intended, whereas misleading is not necessarily

MD-M2: Deception is conscious, whereas misleading is not necessarily

MD-M1 invokes the strongest standard of psychological culpability for deception, namely, intent. Though Carson takes intent to be necessary for deception, my impression is that he is effectively committed to something more like MD-M2.

Carson's remarks suggest that he is attempting to preserve something about the ordinary meaning of the terms: That "deception" is used to signal a negative moral evaluation in a way that "misleading" is not.¹⁰ However, adopting a position like MD-M1 is not the best way to achieve this aim. If we accept that unintentional, reckless behaviour is blameworthy, then it is plausible that reckless behaviour resulting in the formation/reinforcement of a false belief is immoral even in the absence of deceptive intent. Thus, perhaps we should like to use the term bearing negative moral connotations to describe such behaviour (i.e., "deception), however, on MD-M1, only the morally benign term would be apt (i.e., "misleading").

MD-M2 is less restrictive than MD-M1, though it still might be too restrictive to capture our intuitions and practices about moral blameworthiness. On this view, any of intent, knowledge, or recklessness are sufficient to qualify an act as deceptive and not just misleading. The weakest element that qualifies behaviour as deceptive here is recklessness which, in legal terms, amounts to a conscious awareness of risk (Brady, 1990).¹¹ However, many kinds of

¹¹ "Under the Model Penal Code's definition, one is reckless where the agent "consciously disregards" a substantial and unjustifiable risk and negligent where the person "fails to perceive" such a risk.'² In England the Law Commission's Report on the Mental Element in Crime proposed a similar test of recklessness as to result: "Did the person whose conduct is in issue foresee that his conduct might produce the result and if so, was it unreasonable for him to take the risk of producing it?"

behaviour are morally blameworthy even if the agent is unaware of their foreseeable risks. The law accounts for this with the notion of criminal negligence, whereby a person may be held liable for unintentional and unconscious acts if they breach a duty of care. Thus, we may also wish to reserve “deception” to describe those cases where a person’s negligent behaviour results in the formation/reinforcement of a false belief in the mind of another, even if they were not consciously aware of the risk. MD-M2 thwarts this wish.

It may be that Carson does not wish to count negligent cases as deceptive; this is implied by his stating that the negative connotations associated with “deception” may be lost if we choose to label completely unintentional cases as deceptive (Carson, 2010, pp. 47). However, he bases this worry on the assumption that people are often *blameless* for completely unintentional behaviour that causes false beliefs. Even if this assumption is true, as long as there are a significant number of cases that are morally blameworthy despite being unintentional, the intent to deceive alone will be insufficient for the function of distinguishing a morally blameworthy act that causes/reinforces a false belief from a blameless one. And if there are enough *completely* unintentional cases, where a person is morally blameworthy despite being unaware, neither will the conscious element of MD-M2 capture all that is immoral as deception. So, if diverging uses of “deception” and “misleading” are supposed to express the difference between such acts that are blameworthy and the rest that are not, then MD-M1 would be the wrong way to go about explaining the distinction. Given that Carson’s aim is to use the terms in precisely this way, it would make more sense for him to take up a position like MD-M2. However, if we want to admit that people are often morally blameworthy for causing/reinforcing false beliefs even if they are unaware of the risk, then we should reject MD-M2.

Conclusion

So, I have argued that neither intent nor conscious awareness can distinguish an immoral act that causes/reinforces a false belief from a moral one. This does not preclude the possibility that there is a moral distinction between misleading and deception— we may still postulate that

MD-M3: Deception is wrongful, misleading is not necessarily

I will leave aside the question of whether deception is wrongful; recall that we are studying the misleading-deception to help us figure out what misleading really is. However, MD-M3 tells us very little about either of deception or misleading— all it says is that there is *something* that makes the former and not the latter wrong, but not *why*. Since intent and conscious awareness could not function as that something, we are left with no idea of what makes misleading supposedly benign in nature.

We have, however, come a long way in our attempt to characterize misleading. In the first section, we established that misleading and deception are not identical, because deception is always successful in causing or reinforcing a false belief, whereas misleading is not. From this, it also followed that misleading is not a form of deception. In the next chapter, I will attempt to reconcile all that we have established this far into a single, compact definition of misleading. Then, I will return to the topic of the wrongfulness of misleading to address MD-M3.

Chapter 3: A Misleading Tendency

Liars misrepresent their beliefs, and often do so with the goal of deceiving their audience, deceivers cause or reinforce false beliefs in the minds of others, but what do misleaders do? In this chapter, I argue that a misleading communicative act has a tendency to cause people to reason badly. I start by explaining what I mean by ‘tendency’ and ‘cause people to reason badly’. Then, I apply my definition to some paradigm cases of misleading. Finally, I will address the wrongfulness of misleading, and explain how my definition relates misleading to lying and deception.

Misleading: My Definition

I propose to define misleading communicative acts as follows:

Misleading: A communicative act C is misleading iff C has a tendency to cause others to reason badly.

I should break down precisely what this means. On my definition, the essential element of a misleading communicative act is that it possesses a tendency to affect the rational and deliberative processes of those who observe it, such that they are left unable to reason or deliberate well. Qua ‘tendency,’ it is not necessary that a misleading communicative act causes someone to reason badly *in actual fact*, it need only be likely to have that effect; this is similar to how a particular ice cube possesses a melting tendency irrespective of whether it ever melts. Therefore, success is not a necessary condition for misleading on my account.

Further to this, that a communicative act actually causes someone to reason badly is not on its own sufficient to say that it was misleading on my account. To illustrate, imagine a person, A, who only ever utters simple facts that they know are true. A’s friend, B, has spent so much time with A that they have come to believe that all people are this reliable. So, B goes out into

the world trusting that everything anyone says is a simple fact. Even though A's behaviour may have caused B to be so naïve, it would be strange to say that it possessed the tendency to do so. Indeed, uttering facts is generally unlikely to be misleading, and B's behaviour in this case is quite peculiar. Thus, it would be unintuitive to say that A's behaviour had a misleading *tendency*.

The likelihood that a communicative act will cause others to reason badly is determined by a range of contingent contextual facts including, for example, who performs the act, who is present when the act is performed, the social-political climate under which the actor and observers find themselves, and the background assumptions each of them brings to the context. The range of contextual information that is relevant to determining whether an act is misleading is indefinite and perhaps even infinite, which makes it extremely difficult to ascertain whether a given C is misleading. This leaves the boundary of what counts as misleading quite vague; however, I do not take this implication to be a problem. In fact, it is quite intuitive to think that there are many edge cases which may be misleading in some respects and not misleading in others.

One plausible way of figuring out whether a communicative act has a misleading tendency is if a reasonable actor could foresee that their behaviour risks causing someone to reason badly. For instance, if I tell a room of children that nicotine vapes taste like gummy bears, I should be able to foresee the likelihood that some will take it to mean that vaping is a great idea. On accounts of misleading where success is necessary, we could not say that my communicative act here was misleading unless it was the case that the risk manifested for at least one child. On my view, the foreseeability of this risk is evidence for the misleading tendency of the communicative act. This has some interesting moral implications, which I will discuss later. Before that, I will discuss what I take to be the main virtues of my definition.

Misleading without causing false beliefs

In the last chapter, I argued that misleading does not necessarily cause a *false* belief, since people can mislead by implicating something that they believe is false even if it turns out to be true. Another kind of phenomenon that my view accounts for is misleading without causing anything like a belief. Beliefs are propositional, which is to say that they bear truth-conditional content. However, misleading does not always involve the communication of propositional content. For instance, my view accommodates misleading via covert intentional or unintentional dogwhistles. A dogwhistle is a linguistic mechanism designed to bring forth pre-existing attitudes in specific members of the audience. As Saul (2018) explains,

what is dogwhistled is not a particular proposition. Instead, certain pre-existing attitudes are brought to salience. This means that any theory relying on the communication (via semantics or pragmatics) of a particular proposition will fail. (pp. 22)

To use some of Saul's examples, politicians will use dogwhistle terms like 'welfare' or 'inner-city' to trigger subconscious, racist attitudes in voters, thus making them (voters) more likely to let racial bias affect their voting decisions. For a more detailed example, imagine a politician, A, who uses a dogwhistle to show support for anti-Muslim immigration policy:

(12) We need to keep our borders closed to terrorists

A is aware that certain members of their audience hold the belief that 'terrorist' is synonymous with 'muslim,' and that they will therefore interpret the utterance to mean:

(12*) We need to keep our borders closed to Muslims

Thus, without saying anything explicitly racist, A is able to garner voter support from racist voters.

This illustrates how audience members can interpret a secondary proposition from A's use of the dogwhistle. However, Saul's theory about *covert* intentional and unintentional

dogwhistles shows how a communicative act can influence audience members decisions without their ever interpreting or representing any particular proposition. Consider an audience member, B, who hears A utter (12). B does not tacitly believe that Muslims should be kept out, nor does B believe that “terrorist” refers to Muslims. However, B is a part of a complex social system where even people with explicit anti-racist commitments are susceptible to subconscious racist attitudes and implicit biases. Thus, without ever representing a racist proposition, B’s implicit biases are triggered by (12), where they can influence B’s decisions on election day. I take for granted that reasoning affected by either explicit or implicit racial biases amounts to bad reasoning. Thus, my view allows us to say that covert dogwhistling can be misleading even though they do not cause or reinforce any particular false belief.

People can mislead via non-propositional communication using other perlocutionary linguistic devices, like figleaves. A racial figleaf is added to a racist utterance to prevent an audience from recognizing that the utterance was racist. For example, preceding an utterance with “I am not a racist but...” is often an attempt to ‘cover up’ the racist speech which follows (Saul, 2017). In “Racial figleaves, the shifting boundaries of the permissible, and the rise of Donald Trump,” Saul explains how a racial figleaves invoke an ‘inference-blocking mechanism’:

A racial figleaf is, generally speaking, an attempt to block an inference from the fact that the speaker has made an openly racist utterance R to a claim like [C]:

[C] The speaker is racist.

Given the ideology of personalism, this blocking will also have the result that the utterance R itself is no longer seen as racist. (pp. 107)

The inference-blocking effect of the figleaf is perhaps a straightforward count of deception, since the speaker intends to prevent the listener from forming a true belief like C, or the true belief that

the original utterance was racist. However, according to Saul, another effect of racial figleaves is that they shift the boundaries of what is seen as morally acceptable speech.

Consider, for example, the moral obligation that many of us feel to stand-up to racism. Those of us who subscribe to this obligation should feel prompted to object when someone says something explicitly racist. However, as Saul points out:

If a figleaf has been uttered, there is room for doubt about the racism of the utterance, which removes the otherwise uncomfortably present obligation to object to racism.

(pp. 109)

The removal of this obligation thereby makes the utterance, and other utterances with similar contents, morally permissible in the eyes of those who otherwise have non-racist commitments. In addition to inference-blocking, my view makes it possible to say that figleaves are misleading, because they influence audience members' abilities to reason well. In particular, audience members are misled to make bad *deliberative* inferences, as opposed to *epistemic* inferences; the difference being that the former is about deciding on a particular course of action, whereas the latter is about arriving at a particular piece of knowledge. The use of racial figleaves has a tendency to cause its audience members to make bad inferences about how to respond to the racist utterance. If an audience member with anti-racist commitments does not recognize the pernicious use of the figleaf, they will be prompted to overlook racism where they would otherwise have objected to it. My account lets us say that the speaker's use of the figleaf was misleading in addition to its being inference-blocking.

Misleading without success

Being solely concerned with false beliefs, prior accounts of misleading overlooked the possibility that people could be misled without ever representing something propositional in form. My definition of misleading helps us to identify the non-propositional effect that

mechanisms like dogwhistles and figleaves have on audience members' abilities to reason well. Another virtue of this definition is that, unlike deception, misleading speech and behaviours do not need to be successful. The case where I told children that nicotine vapes taste like gummy bears is one example of a case that should count as misleading irrespective of whether it succeeds at making anyone reason badly. Consider another case where there is neither an intent nor a resulting false belief:

A advertises herself as a specialist nutrologist when no formal nutrology qualifications exist in Canada, without thinking about how the term 'specialist' could imply that she does. None of her clients believe that she has a formal qualification from Canada; nor do they concern themselves with the matter.

Even though no one has in fact been led to believe the false proposition that the nutrologist has formal qualifications in Canada, it is easy to see how her advertising herself as a specialist could suggest that she does. Indeed, her advertisement seems to meet the criteria for misleading commercial speech, since it is very reasonable to interpret the term 'specialist' as meaning that someone does possess a formal qualification. My definition makes it possible to say that, regardless of her intent and the resulting mental states of her clients, the nutrologist's advertisement is nevertheless misleading.

Objection: Isn't everything misleading?

Earlier, I raised a limitation regarding our ability to ascertain whether a communicative act is misleading. That is because whether a communicative act has a misleading tendency depends on a broad and indefinite range of contextual factors, which we have limited knowledge about. The cases of the gummy bear vapes and the Canadian nutrologist are engineered to grant

us privileged access to contextual information such as the ages and mental states of audience members. However, real life cases of misleading are seldom so clear— we do not always know who our audience will be, such is the case with speech online or in front of a large crowd. Nor are we omniscient of the many social and political forces which may influence our communicative practices. Thus, there will be no clear formula for dividing up misleading from non-misleading except for our best judgment about how likely the communicative act is to make people reason poorly.

However, that my definition generates vague boundaries for what counts as misleading should not be seen as a drawback— in reality, misleading is a messy phenomenon, and we will often have conflicting intuitions about whether a given case was misleading or not. This is no less true for lying and deception. If a person asserts a proposition which they neither believe nor disbelieve, perhaps they just have a hunch that it could be false but aren't sure, it would be difficult for us to tell with certainty whether they lied. If a person causes someone to believe something false, but no one on earth knows that it is false yet (e.g., if I convince you to believe that a leaf will drop onto that bench at 2:00pm tomorrow), can we be sure about whether deception has taken place? Communication in the real world is messy, and so misleading is doomed to be.

Vagueness aside, a possible problem with my definition is that it is too broad. After all, it may be that any communicative act has a tendency to make someone reason badly given the right set of circumstances. Other views of misleading were at least able to narrow down to only those cases that resulted in false beliefs but, by eliminating this constraining element, it turns out that misleading is far more ubiquitous on my account. I also do not find this to be much of a problem. I think most would agree that misleading is far more common than lying and deception; the

number of people who lie or purposely inculcate false beliefs in the minds of their audiences is surely less than the number of people who inadvertently say or do things that are likely to arouse confusion. That said, the misleading character of a communicative act is likely to admit of degrees, such that one may be more or less misleading than another depending on how likely it is to cause someone to reason badly. For instance, recall the example where B was caused to reason badly by A's austere commitment to uttering simple facts. A's behaviour may well possess a slight misleading tendency, if not due to the low threshold for what might trigger an audience as naïve as B into reasoning poorly. Still, uttering truths is sure to have a much weaker misleading tendency than uttering lies, which are inherently misleading given that they altogether block the audience from inferring the liar's genuine mental states. Thus, while many if not most communicative acts may possess a weak tendency to mislead, the stronger and more obvious cases of misleading are less common.

From a practical standpoint, there are ways of measuring the misleading tendency of a communicative act. For instance, if one is speaking to an audience that is already disposed to reason poorly, then the chance that their speech will mislead is probably higher than if they were speaking to an audience of careful and critical minded thinkers. For example, imagine trying to communicate to an audience of misogynists. Being misogynists, they are disposed to reason poorly whenever gender is a salient topic. Thus, talking about gender around them will have a much stronger misleading tendency than talking about gender with a group of egalitarians. Another example would be communicating in ways that are unclear, such as someone who communicates with an intense sort of sarcasm that people find difficult to interpret and understand. Communicating in this way has a stronger tendency to confuse others than

communicating clearly and concisely; though it may be more entertaining, blunt sarcasm is far more likely to mislead than direct speech.

What's wrong with misleading?

Earlier in this chapter, I suggested that reasonable foreseeability could serve as evidence for the misleading tendency of a communicative act; that if a person could have foreseen the risk that their behaviour would cause others to reason poorly, then the conclusion that the act was misleading would be justified. Moreover, if we agree that it is immoral to recklessly cause others to reason poorly, reasonable foreseeability should be sufficient to establish that the person is morally blameworthy for the act. However, this last point raises the questions: Is it really wrong to make others reason poorly? To what extent are we morally blameworthy for communicative acts that have a tendency to mislead?

One way that philosophers have tried to explain the immorality of acts like lying, misleading, and deception is by appeal to their unethical consequences. Thus, it could be said that successful lying and deception are wrong because they cause false beliefs in the minds of others. Even if this is a plausible view with respect to lying and deception, it cannot be for my account of misleading, since causing false beliefs is not an essential part of the picture; the potentially harmful consequences of a misleading tendency are not always manifest. So, something other than consequences must explain the immorality of misleading.

Alternatively, the wrong of lying and deception can be explained by the psychological states of the liar/deceiver— if they intended to misrepresent their beliefs or cause false beliefs in the minds of others, then this might be enough to establish moral blameworthiness. However, this strategy also could not explain why misleading is wrong, since my definition does not say anything about the mental states of the person performing the communicative act. That a

misleader is intentful, knowing, reckless or negligent might help us to justify the judgement that their act was misleading, but nothing akin to *mens rea* is necessary on my account.

In the last chapter, we discussed Carson's theory that the term "deception" carries with it a negative moral connotation whereas "misleading" does not, which he explained in terms of intent: Deception is necessarily intentful and therefore immoral, whereas misleading is not. I argued that intent could not function to distinguish moral from immoral behaviour resulting in false beliefs. It turned out that intent was too strong a standard, since it left out cases where someone knowingly, recklessly or negligently caused someone to believe something false— all of which we would intuitively regard as immoral. I ended the chapter by raising the possibility that deception is the wrongful counterpart to misleading:

MD-M3: Deception is wrongful, misleading is not necessarily

This would require anything from intent to negligence for a false belief-causing act to be aptly labelled as deception— otherwise, the act should be called misleading. I think that, if one insists on maintaining that the term "deception" has a stronger connotation of blame than "misleading," MD-M3 would be the most plausible way of explaining why. As a claim about language use, I do not actually know whether it is empirically true, though, I suspect that it is not. My impression is that "misleading" and "deception" only differ in the *kind* of negative moral connotations they bear. But MD-M3 is not phrased as a claim about language, it is a way of distinguishing the phenomena themselves and, as such, I think it is mistaken for at least two reasons: Not all deception is wrongful; and misleading, on my account, is wrongful most of the time.

It is relatively easy to think of cases of deception that are not wrongful, even if they are intended. For instance, imagine a woman who wishes to have an abortion but knows that, if her

partner finds out she is pregnant, they will try to coerce her out of it. To preserve her privacy and right to choose an abortion, she deceives her partner into believing that she is not pregnant. Here, the woman's rights to privacy and choice should override whatever duty she might have to not cause false beliefs in the minds of others. Many other cases where privacy is concerned are like this; for instance, an employer who asks illegally for personal medical records or biometric data. But, the presence of an overriding duty or right is not the only circumstance under which we would find deception permissible. Imagine your partner asks if those pants make her butt look big, or if your well-intending father asks if you liked his creative rendition of black pudding. Even if deception is wrong most of the time, cases like these suggest that it is not necessarily wrong to cause someone to believe something false.

On the other hand, I cannot think of a counter-case where it would be okay to cause someone to reason poorly. Where intent, recklessness, and negligence are present, causing someone to reason badly is surely immoral. Doing so inadvertently may not be strictly blameworthy, but there doesn't seem to be a case where we would be comfortable with this outcome in the way that we would be in the pregnant woman's example, or when you deceived your partner and your dad. Causing people to reason poorly is always a bad thing.

However, as I said a few paragraphs ago, this consequence does not suggest that misleading is inherently blameworthy, since a communicative act need only have a *tendency* to cause bad reasoning in order to count as misleading; the risk is not always realized. However, even if an agent is not psychologically blameworthy for a totally inadvertent misleading act, I do actually think that there is something inherently wrong about it. This is analogous to how, for example, a person who forgets to wear a mask during a pandemic is doing something wrong, even if they totally lack any element of intent. Or, to use an even less contentious example, we

would think it is bad for someone to drive a car under the influence of a drug which causes drowsiness, even if they had no idea about the side-effect. Other examples might be taking a job on the basis of white privilege, or accidentally misgendering someone. Sometimes, the risk that an act poses to society as a whole is in itself a reason to think that the act is wrong, irrespective of the mental states of any particular agent.

Attas (1999)'s "What's Wrong with 'Deceptive' Advertising?" provides a framework that will be useful to elaborate this point in the case of misleading. Attas writes:

Rather than looking at the consequences of a single *act* of a misleading *advertisement*, I suggest we look at the consequences of the *practise* of misleading *advertising*. That is, the situation where no disincentive is legally instituted for negligence in the distortive transmission of information in advertising. This also suggests that instead of looking at the consequences of *a* consumer holding a false belief due to a misleading advertising message, we look at the consequence of a high likelihood of consumers in general being misled. Thus, the morally relevant harm is not necessarily or uniquely that imposed on the misled person, but the harm imposed on society at large. (pp. 56)

On Attas's account, the wrong of misleading advertising lies not in its actual consequences, but rather in the consequences that arise when advertisements with a misleading tendency are allowed to proliferate in society.

As a final example, recall the case of the Canadian nutrologist: Although she did not intend to deceive anyone, and did not actually cause anyone to believe anything false, we can isolate the wrongfulness of her actions by considering how society would turn out if such advertisements were broadcast on TV or online. On my definition, misleading entails a probabilistic tendency to cause people to reason badly. Given that many kinds of communicative

acts (including advertisements) today are seen by audiences of thousands online or on TV, the risk that someone will be triggered to reason poorly is highly likely to be realized. That members of a society are constantly exposed to misleading communication is itself a threat to that society's ability to flourish. Thus, we should maintain that behaviors which have a misleading tendency are inherently wrongful if only for the risk that that they pose to society at large.

What does this mean for the moral blameworthiness of an agent who misleads? Well, while it might be too strong to say that people ought never perform a communicative act which has a misleading tendency, given the difficulties with being able to identify when one's behaviours will in fact have this tendency, it might still be reasonable to insist that people ought to be careful about the communicative acts they do perform, and more thoughtful about how they might affect others. Someone who does not care whether they are misleading, or never thinks about whether their communicative acts could have negative epistemic consequences for those around them is likely to be contributing to a potentially harmful social condition where important values like trust and cooperation are at stake. I take this to be an inherent wrong associated with misleading.

Lying, deception, and a misleading tendency

If I am right about this last point, then MD-M3 turns out to be a rather unattractive way of explaining the difference between misleading and deception; if lots of deception *is not* wrongful, and lots of misleading *is* wrongful, then distinguishing the two on moral grounds is probably the wrong way to go about it. This brings me to my final endeavor, which is to consider how my definition of misleading relates it to the other concepts we have looked at in detail here, lying and deception. This happens to be quite straightforward.

My account is compatible with MD-S2, which stated that deception is a form of misleading. I take it that causing someone to believe something false is just one way of causing them to reason poorly. Moreover, on my account, lying is a form of misleading. Even if lying is not successful, saying something believed to be false will always cause people to reason poorly about the mental states of the speaker. Shiffrin (2019) takes this to be the wrong of lying:

I take the wrong of the lie to be grounded in the liar's abuse of the channel of testimony as a reliable and distinctive mode of access to the contents of the speaker's mind. (pp. 72)

So, on my view, lying and deceiving are both ways of misleading; that is, they are communicative acts which have a tendency to cause others to reason poorly.

As Saul argued, the difference between lying and merely misleading is the difference between saying something false and communicating something false without literally saying it. This is not in tension with lying being a form of misleading, it just describes what it means to mislead without lying. Stokke's description of merely misleading is more specific in calling it non-lying *deception*, since causing a false belief is deceptive on all the accounts we looked at here. In any case, my view implies that all deception is misleading in nature, as well as all lies, in a way that is compatible with both Saul's and Stokke's views. The original contribution my account makes to the literature on the lying-misleading and misleading-deception distinctions is in proposing that misleading is the overarching concept. Lying and deception share a tendency to make others reason badly and, as such, they are forms of misleading.

Conclusion

This thesis is my proposal to redefine how philosophers think and write about misleading. Until now, the literature on misleading and nearby concepts like lying and deception has been heavily focused on two elements: The intent to deceive, and the causing of false beliefs. However, lots of what we would intuitively label as misleading involves neither of these elements. Moreover, as we saw in Chapter 2, laws about misleading advertising typically do not require proof of intent or proof that someone came to believe something false. Another problem entrenched in the literature is that philosophers tend to conflate misleading with deception, or they presume that misleading is just inadvertent or morally benign deception. In Chapter 2, I argued that both views are mistaken. Finally, in Chapter 3, I defended my definition of misleading: A communicative act *C* is misleading iff *C* has a tendency to cause others to reason badly. The hallmark feature of this definition is the misleading ‘tendency,’ which accounts for cases where behaviour is likely to cause people to make reason badly even if it does not in actual fact. Moreover, my definition lets us characterize communicative acts which have non-propositional effects on their audience’s minds, such as dogwhistles and figleaves. Lastly, I explained that the moral wrong of misleading on my view does not turn on either of psychological culpability or deceptive consequences; rather, misleading is wrongful insofar as the proliferation of misleading behaviour has a negative impact on society as a whole.

References

- Attas, D. (1999). What's Wrong with "Deceptive" Advertising? *Journal of Business Ethics*, 21(1), 49–59. <https://doi.org/10.1023/A:1005985723779>
- Augustine, S. (1952). "Lying". *Treatises on Various Subjects (The Fathers of the Church, Volume 16)* (R. J. Deferrari, Ed.; Sister Mary Sarah Muldowney, H. B. Jaffee, Sister Mary Francis McDonald, Sister Luanne Meagher, Sister M. Clement Eagan, & M. E. DeFerrari, Trans.). Catholic University of America Press. <https://doi.org/10.2307/j.ctt32b2mf>
- Bates v. State Bar of Ariz., 433 US 350 (Supreme Court 1977).
- Bi, K. (2015). What Is "False or Misleading" Off-Label Promotion? *The University of Chicago Law Review*, 82(2), 975–1021.
- Beauvoir S & de Kirkup J. (1959). *Memoirs of a dutiful daughter* [First edition]. World Publishing Company.
- Camp, E. (2018). Insinuation, common ground. *New Work on Speech Acts*, 40.
- Carson, T. L. (2010). *Lying and deception: Theory and practice*. OUP Oxford.
- Donnellan, K. S. (1966). Reference and Definite Descriptions. *The Philosophical Review*, 75(3), 281–304. <https://doi.org/10.2307/2183143>
- Fallis, D. (2018). What Is Deceptive Lying? *Lying: Language, Knowledge, Ethics, Politics*.
- Grice, H. P. (1975). Logic and conversation. In *Speech acts* (pp. 41–58). Brill.
- Matter referred on 21 April 2022 (conduct of Rt Hon Boris Johnson): Final Report*. (n.d.).
- McKittrick, J. (Jennifer L. (1999). *The metaphysics of dispositions* [Thesis, Massachusetts Institute of Technology]. <https://dspace.mit.edu/handle/1721.1/9368>

- Michaelson, E., Sterken, R., & Pepp, J. (2019). What's New About Fake News? *Journal of Ethics and Social Philosophy*, 16(2), Article 2. <https://doi.org/10.26556/jesp.v16i2.629>
- Santana, C. (2022). What's wrong with dogwhistles. *Journal of Social Philosophy*, 53(3), 387–403.
- Saul, J. (2012). Just go ahead and lie. *Analysis*, 72(1), 3–9.
- Saul, J. (2018a). Dogwhistles, political manipulation, and philosophy of language. *New Work on Speech Acts*, 360, 84.
- Saul, J. (2018b). Negligent Falsehood, White Ignorance, and False News. In E. Michaelson & A. Stokke (Eds.), *Lying: Language, Knowledge, Ethics, and Politics* (p. 0). Oxford University Press. <https://doi.org/10.1093/oso/9780198743965.003.0013>
- Saul, J. M. (2012). *Lying, misleading, and what is said: An exploration in philosophy of language and in ethics*. Oxford University Press.
- Saul, J. M. (2017). Racial Figleaves, the Shifting Boundaries of the Permissible, and the Rise of Donald Trump. *Philosophical Topics*, 45(2), 97–116.
- Shiffrin, S. V. (2019). I—Learning about Deception from Lawyers. *Aristotelian Society Supplementary Volume*, 93(1), 69–90. <https://doi.org/10.1093/arisup/akz010>
- Shiffrin, S. V. (2021). Back-door Lies and Promising under Coercion. *Mind*, 130(517), 265–279. <https://doi.org/10.1093/mind/fzaa039>
- Stainton, R. (2006). *Words and thoughts: Subsentences, ellipsis, and the philosophy of language*. Oxford University Press.
- Stein, Howard. (1992). “Was Carnap Entirely Wrong, after All?”, *Synthese*, 93(1–2): 275–295. doi:10.1007/BF00869429

Stokke, A. (2013). Lying, Deceiving, and Misleading. *Philosophy Compass*, 8(4), 348–359.

<https://doi.org/10.1111/phc3.12022>