

TALKING TOGETHER:
ON THE NORMATIVE SIGNIFICANCE OF
OUR DECLARATIVE PRACTICES

A Dissertation

Presented to the Faculty of the Graduate School
of Cornell University

In Partial Fulfillment of the Requirements for the Degree of
Doctor of Philosophy

By

Alejandro Vesga

August 2023

©2023 Alejandro Vesga

TALKING TOGETHER:
ON THE NORMATIVE SIGNIFICANCE
OF OUR DECLARATIVE PRACTICES

Alejandro Vesga, Ph. D.

Cornell University

Abstract: This dissertation places cooperative action at the core of a systematic account of the linguistic, epistemic, ethical, and social dimensions of our declarative practices—of the ways we communicate to each other about the world. Say you and I are talking about that time Sam was detained and suddenly released by the police, and I assert “He bribed them.” I begin by arguing that my act is best understood as the act of cooperatively offering information to achieve the shared goals that structure our conversation. This means that, if I was only guessing and Sam did not actually bribe the police, my failure is an instrumental failure to achieve what I committed to do with you—just as passing you baking soda instead of salt is a failure in our joint project of making pasta. Crucially, unlike other views in the literature, this view neatly explains the distinctively second-personal entitlement you have to reproach me for my failure; and it avoids the problems of accounts that explain my failure in terms of the evolutionary function of assertion. Moreover, the view I develop also explains why falsely insinuating that Sam bribed the police by merely saying “he offered them hundreds of reasons to let him go” is often taken to be less committal, and thus less morally wrong, than my assertion above. The normative distinction between committal and non-committal speech is best captured by the distinction between cooperative and uncooperative speech—as opposed to the distinction between literal and non-literal speech, as is widely assumed in the literature. And, finally, I argue that the relationship between knowledge and successful action gives us reason to think that assertion requires something at least close to knowing what we say, and not, for instance, rationally believing it, or merely saying true things.

BIOGRAPHICAL SKETCH

Alejandro Vesga was born in Bogota, Colombia. He did his BA and MA in philosophy at the Universidad de los Andes, Bogotá, Colombia. Alejandro started studying in the Sage School of Philosophy at Cornell University in 2017. In 2023, he graduated with a PhD in Philosophy and a graduate minor in Cognitive Science.

To my family, blood and not.

ACKNOWLEDGEMENTS

I am often reluctant to say that I am proud of other people's accomplishments because I want to avoid the connotation that I had anything to do with, or deserve any merit for, what they achieved. By that logic, I would be elated if my work made the people mentioned here at least somewhat proud. For I am more than happy to admit that they are part of this dissertation; after all, dare I say, it is the outcome of what we have done together.

First, I have my advisors to thank. One of the aspects of this dissertation that makes me especially proud is to see materialized in what I have to say about communication all that I have learned from Carlotta, Starr, Shaun, and Kate about action, language, epistemology, normativity, and cognition.

Carlotta's expertise, patience, and brilliant and incisive feedback have been crucial to bringing this dissertation together. Before and throughout the beginning, middle, and end of a global pandemic, Carlotta has met with me to talk about incessant drafts, expanded and corrected my ideas when they needed a push, and gave me invaluable professional advice. Her generous expertise and skills on language, action, and epistemology were essential in developing my own. Her advising lead and trust in my own ideas was invaluable to me. Likewise, I am grateful to Starr, not least because they believed in my project before I did. From the beginning of my PhD, they encouraged me to pursue what I felt was just a confused idea. I have learned much from their work and their way of thinking about the interface between meaning, communication, and sociality.

I am grateful to Kate for generously helping me shape my initially obscure intuition that the reasons I care so much about communication are deeply connected to the reasons I care about social injustice and the normatively grounding role of other people in our moral lives. She trusted that I had something to say about this and gave invaluable feedback on work that was out of her wheelhouse and, finally, on work that was not. And to Shaun, I am grateful for being open to seriously discussing themes and fields I was truly ignorant about. Our conversations and work on social and normative cognition illuminated an interdisciplinary dimension of my research and interest that I was not expecting when I came to Cornell. Our non-dissertation-related work has been immensely satisfying and I am proud to see how it has shaped my current and future research.

I was not only fortunate because I had four philosophical powerhouses in my committee; but because they were kind and supportive about my work, especially in moments when I failed to be so with myself. I do regret bringing Morten Christiansen to the committee later than I should have. His work and our conversations have helped me strengthen and sharpen my arguments about the normativity of communication. I thank him for his willingness to step in. Hopefully, this will not be the end of our conversations on language and communication.

I have to thank the Cornell faculty who met with me to discuss my thesis and other related, and admittedly disparate, ideas. My conversations with Michelle Kosch and Julia Markovits had a vital role in shaping my understanding of the ethical and normative import of shared agency. I also have to thank Nico Silins, Arc Kocurek, and Dave Shoemaker, for patiently listening to half-baked ideas and answering wisely to the subsequent question ‘does this make sense to you?’. I am grateful to Sarah Murray in linguistics and David Pizarro in psychology for listening to my ideas and providing guidance on where to look next outside of philosophy. And although she is not at Cornell anymore, I am thankful to Tamar Kushnir, whose class on the psychology of imagination and advice put me on an interdisciplinary path that has been enormously fruitful.

Beyond Cornell, I am awed by the list of people who have helped me become an at least decent philosopher. My MA advisors, Tomás Barrero and Santiago Amaya, have been a constant source of academic and professional support and advice throughout the PhD. This makes me particularly happy, as the inception of this dissertation’s theme was developed in my MA thesis and under their tutelage. I am also grateful to Clayton Littlejohn, whose advice made my stay at King’s College productive and exciting, and to Peter Graham, who let me to sit in on his class on testimony and has been an invaluable interlocutor since then. My dissertation and philosophical acumen have also been greatly improved thanks to my conversations with David Papineau, Liz Camp, Jenny Saul, Lizzie Fricker, Luca Ferrero, Manuel Vargas, Jennifer Morton, José Andrés Forero, David Rey, Julienne Dutant, and Elliot Christiansen.

If I can leave one piece of advice on the page, it is that people, not ideas, make or break a PhD. My philosopher friends kept me personally and philosophically sane. I am proud to have grown academically and professionally in tandem with Camilo Martinez and to see how our ideas have grown together too. I am enormously grateful to Libby Southgate, Maria Camila Castro, Migdalia Arcila, Fran Fairbairn, Geoff Weiss, Bobbi Cohn, Zeyad Elnabosy,

Alicia Patterson, Vikram Kumar, Hanna Winker-Olick, and Brianna Zgurich for their company and support.

The Colombian community was also essential for survival. To Diana Obregón, who I met at the airport on our first day in Ithaca and has been present ever since, I am eternally grateful—not least for bringing David and Amalia to my life. I am grateful for the kindness and friendship of Maria Camila Castro, Migdalia Arcila (yes, again), Isabel Calderón, Roberto Ibañez (honorary Colombian), Diana Fajardo, Ana María Porras, Juan José Silva, Eugenio Valdéz, and the many others that made my life a good life in Ithaca.

And to my family, the greatest gratefulness. I hope to be able to return and give others at least a small parcel of the support, company, and happiness that my mom, my dad, and my two brothers have given me. I am immensely fortunate to have them as a home to come back. But also, I am grateful for the friends I have had for at least 27 years and who have become my family as well. I really (really) wish there was a better name to denote them, but to the Tilamigos I am forever grateful and proud to see them in what I do; for answering every call, having me in their homes, and forgiving my absence, *gracias*.

Por último, dejo una nota en español para reconocer el valor de la comunidad lingüística que hace de mí lo que soy, incluso escribiendo en otro idioma y en otro país. Espero volver, sabio y mejor, a hacer filosofía en el idioma que me crio; y que el malestar cultural que tengo se traduzca en una filosofía americana.

Como dice mi mamá, a todos gracias, gracias, gracias.

TABLE OF CONTENTS

INTRODUCTION	1
CHAPTER 1 Communicating Testimonial Commitment	12
CHAPTER 2 The Ethics of Communicative Commitment	34
CHAPTER 3 The Dual Nature of Testimony	66
CHAPTER 4 Can the Norm of Assertion be Functional?	88
CHAPTER 5 Talking with Others: The Cooperative Norm of assertion	100

CHAPTER ABSTRACTS

Chapter 1: Communicating Testimonial Commitment. I argue for the Cooperative Warrant Thesis (CWT), according to which the determinants of testimonial contents in communication are given by the practical requirements of cooperative action. This thesis distances itself from conventionalist views, according to which testimony must be strictly bounded by conventions of speech. CWT proves explanatorily better than conventionalism on several accounts. It offers a principled and accurate criterion to distinguish between testimonial and non-testimonial communication. In being goal-sensitive, this criterion captures the role of weak and robust cooperation in determining the contents to which speakers testify or fail to testify. And, finally, it yields a principled explanation of why testimony entails the epistemic commitments that distinguish it as an epistemic source.

Chapter 2: The Ethics of Communicative Commitment. If I were to deceive you into falsely believing that Ana bribed the police, saying ‘she bribed them’ seems worse than merely saying ‘she solved it like a businesswoman.’ This suggests that there is a normative difference between declarative speech acts that commit speakers to a content and declarative speech acts that do not. I argue that a vastly popular rendering of this normative differential is misguided. Call it Literalism: the idea that the normative distinction between committal and noncommittal speech traces the moral distinction between literal and nonliteral speech. Against Literalism, I suggest that there is a pervasive kind of nonliteral speech that systematically carries the same commitments carried by literal speech: Nonliteral cooperative signals. Hence, I argue, we should endorse Cooperativism, which holds that the normative difference between committal and noncommittal speech is best explained by appealing to the distinction between contents that are required and contents that are merely allowed as cooperative moves in a given conversation. Cooperativism, in turn, i) accommodates context-dependent cases of literal and nonliteral commitment, ii) better captures the moral principles that seemed to lend support to Literalism, and iii) offers a principled pathway to theorize about the role of social phenomena (like power, institutions, and social norms) in the ethics of speech.

Chapter 3: The Dual Nature of Testimony. An adequate theory of testimony must not only i) explain why testimony justifies belief and ii) explain why testimony entails second-personal relations of accountability. It shall also answer iii) why these two facts come together; it should, that is, account for a common explanatory link between them. Call it the Dual Nature Desideratum. This desideratum, I argue, reframes the challenges that traditional evidentialist and assurance views in the epistemology of testimony pose to each other. Correspondingly, here I defend a distinctive account of testimony that answers (i)-(iii). Cooperative action, I contend, offers a unified grounding source of the second-personal and evidential natures of

testimony. Cooperative engagement between speaker and addressee enables our testimonial practices. In turn, the practical pressures of cooperative action provide distinctive epistemic grounds that justify belief and entitle audiences to hold speakers accountable for failing to fulfill their cooperative commitments.

Chapter 4: Can the Norm of Assertion be Functional? Several epistemologists hold that there is a constitutive link between the norm of assertion and the speech act of assertion. I present a problem for a novel way of specifying this view, namely etiological functionalism about assertion. Whether the speech act of assertion has the etiological function of generating this or that epistemic good is a contingent matter, determined by the specific selection pressures the practice was selected by. Therefore, etiological functional norms cannot fulfill the constitutive role that epistemologists interested in the norm of assertion have in mind. I offer an alternative path towards a constitutivist functionalism. Functionalists should abandon etiological functions and endorse artifactual functions to account for the constitutive norm of assertion. Such an artifactual functionalist view provides a distinctive path to derive the content norm of assertion, no matter the possible etiological stories of the means by which speakers are able to assert in this or other worlds.

Chapter 5: Talking with Others: The Cooperative Norm of Assertion. Constitutivism about assertion holds that there is a constitutive link between the norm of assertion and the very nature of the speech act of assertion. This family of views has been mostly concerned with what I call here the normative question about assertion: should speakers, for instance, know what they assert, should they have justified belief, or should they say true things? Here I suggest a novel path to answer this question by way of addressing another significantly less visited question: what kind of norm is the norm of assertion? Call this the metanormative question about assertion. I argue that the constitutive norm of assertion is not game-like, social, conventional, or functional. Rather, it is an instrumental norm, stemming from the normativity of cooperative action. This view neatly accommodates the non-constitutive role of other communal norms in the practice of assertion. In turn, this cooperativist view provides a distinctive path to specify the epistemic profile of the norm of assertion. The link between intentional action and success suggests that the norm of assertion requires asserted contents to be believed, true, justified, and safe. Whether this profile amounts to knowledge, however, will depend on your epistemological allegiances.

INTRODUCTION

This dissertation revolves around the speech act, core to our communicative practices, of actively informing others about the world. This same speech act has been of central interest in epistemology, philosophy of language, ethics, and social philosophy. However, it has received many names, and its labeling is a contentious matter. For expository purposes, let us call it assertion in this introduction; but I will have more to say about this choice below.

This act is often (but, crucially, not always) instanced by the use of declarative sentences in serious conversation and is stronger than merely guessing, speculating, or insinuating that something is the case. Take a paradigmatic example. Augusto and Lucía are conversing, and Lucía says, “the sweater is in the suitcase.” Under the right circumstances, Lucía used her words to inform Augusto that the sweater is in the suitcase. This act, we will see, is distinctive insofar as it has an exceptionally rich normative profile. There are, that is, various communicative, epistemic, moral, and social dimensions along which this act can be judged as appropriate or inappropriate—each of which has been the center of attention in different pockets of philosophical interest.

My wager throughout this dissertation is that we can have a unified account of this distinctive speech act and its manifold of normative dimensions by appealing to the instrumental rationality of cooperative action. Assertion, say I, is the act of cooperatively offering information aimed at advancing the shared goals that structure our conversations. And it is in virtue of the rationality of cooperatively doing things together that this single act exhibits the slew of its normative dimensions.

Those who have seen intro to philosophy of language know this is an old idea. It lies at the core of Grice’s widely influential Cooperative Principle. And, at least to my mind, it is also central to Lewis’s understanding of what languages are. Yet, the novelty of what I have to say comes in two prongs. First, the substance of my arguments stems from focusing on an action-theory perspective of non-language-specific cooperative action. The instrumental normativity of cooperative action—of the things we do together—serves as the explanatory fulcrum to provide an account of assertion and its normative profile. Second, if my arguments hold, cooperation is essential to understand a wide range of normative phenomena, usually taken to be out of the purview of philosophy of language. Cooperation is essential not only to individuate the communicative content of assertions but to understand the unity of its epistemic, moral, and social import. Let me say more about these prongs.

Normative Unity

Throughout the dissertation, I explore how the nature of cooperative action explains the various normative aspects of the assertion. Although each of the essays deals with more than one of assertion's normative dimensions, it will be helpful to disentangle and describe each of them separately.

We can ask, for starters, about the *epistemic normativity* of assertion. If things go right, Rafael can learn from Ana's act. He can acquire justified belief, even knowledge, that the sweater is in the suitcase. We can ask, then, what in the world makes it the case that, or in virtue of what, assertion justifies belief? How does assertion, for instance, compare with other epistemic sources, like perception or memory? By the same token, if Lucía lacks the epistemic privilege she presented herself as having—if, say, the sweater is not there or if she was merely guessing—her act would be clearly defective. The obvious question is, what counts as a defective assertion? Answering this question involves explaining what epistemic privilege speakers must have to assert properly. Should they know what they say, justifiably believe, or simply say true things? Or, in terms now classic in epistemology, what is the norm of assertion? Further, a full account of assertion should explain what kind of failure is at stake. Is Ana, for instance, breaching a social norm or convention, or does her assertion simply deviate from the function it was selected for? Put otherwise, what kind of norm is the norm of assertion?

Besides its epistemic dimension, one may also ask about the *communicative normativity* of assertion. To assert appropriately, one must use the appropriate communicative mechanisms. Speakers, for instance, often communicate information in less committal ways that are different from the speech act I am interested in here. Say Pablo and Juan are discussing whether Sergio broke up with his girlfriend, and Juan says “well, he’s going a lot to the gym lately”—to merely insinuate that they did break up. At least intuitively, his act is less committal and altogether different from simply saying, “they broke up.” To assert, then, speakers must communicate appropriately. Put otherwise, there are communicative conditions of assertion. One may ask, then, how should speakers communicate to appropriately assert? Or, under what conditions does an utterance count as an assertion?

This leads to questions about the *moral normativity* of assertion. Say that in the scenario above, Juan is aiming to deceive Pablo into falsely believing that Sergio broke up with his girlfriend. Directly asserting “They broke up” seems worse than simply saying “well, he’s going a lot to the gym lately”. If so, an account of assertion should explain, on the one hand, why intentionally communicating falsehoods constitutes a wrong (if it does at all) and what kind of wrong is involved in deception. And, on the other hand, such an account should explain

why deceptive assertions are normatively distinct from deceptive insinuations (if they are at all). Furthermore, say Juan is telling the truth and is well positioned to know what he is talking about. If Pablo dismisses him for the wrong reasons—say because of Pablo’s prejudices against Juan’s ethnicity—this amounts to an epistemic injustice. Hence, one may ask whether, and if so why, this is a distinctive kind of wrong.

Finally—and this is a less navigated question in the philosophical treatment of assertion—one may ask about the *social normativity* of this act. Assertion, like many other human acts, is social. It is constituted and enabled by mutually responsive interaction between agents. These kinds of human interaction are often modulated and mediated by large-scale social phenomena that go beyond the mere individual beliefs and intentions of the interacting agents. In particular, power relations, institutions, and social norms are crucial analytical categories often used to understand how large-scale sociality constrains and determines our individual and second-personal acts. We may ask, then, how (if at all) do social power, norms, and institutions constrain this speech act, so central to our social lives?

A complete theory of assertion as a unified practice has at least two tasks with respect to its various normative dimensions. First, it must answer these normative questions and provide a valid account of the various norms that govern assertion; what is communicatively, epistemically, morally, and socially required of asserting speakers? Second, it ought to provide a metanormative account of why, or in virtue of what, these norms come together; why is it that this single speech act is subject to such a diverse set of norms? Or, put otherwise, what in the world makes it the case that our assertions have this rich normative profile? As a whole, this dissertation is an attempt to advance (even if not exhaustively) a unified view of assertion that fulfills these two tasks.

Cooperative Communication

Your commitment to the individual goal of planting a tree in your garden today demands some things of you: You should have a tree available, the right tools, and dig a good hole, among other things. It also prohibits certain things: bringing a straw to dig the hole, leaving the garden to go camping, and setting your garden on fire. We can say that the space of these prescribed and proscribed actions—and everything in between—is given by the instrumental normativity of intentional action. Insofar as you are committed to a given goal, this goal structures your agency in terms of what you should and should not do to achieve it.

Crucially for my purposes, humans are not only good at doing complex things. We are exceptionally good at doing things *together*. We often—where ‘often’ feels like a radical

understatement—act toward shared goals in ways that can only be achieved by mutually and responsively interlocking our actions.

The cooperative commitments to advance a shared goal with another agent involve a distinctive set of instrumental requirements. Our commitment to plant a tree together requires us to have a tree, the right tools, and dig a good hole. It also requires us to coordinate in ways that do not dissolve this commitment. If we are really acting together to plant a tree, we must act in ways that preserve and ensure our mutually responsive behavior toward this goal. If I let you plant the tree while I watch from the porch, we are not planting the tree together—even if the world ends up looking how we both wanted. Similarly, if I get in my car and leave, I must make clear to you that I am going for fertilizer on pains of you concluding that our shared goal has been dissolved—even if my own intentions were to plant a tree with you.

In turn, the mutually responsive instrumental structure of cooperative action enables a distinctive kind of inference about other people's behavior. Under the assumption that we are both committed to planting the tree, you can understand my going to the shed as my going for the shovel to dig the hole. This is not only because I am committed to planting the tree, but because I am committed to doing so *with you*. Had I gone to rest, I would have made it clear to you. The mutual acknowledgment of cooperative commitments serves as a distinctive kind of warrant for our beliefs about other people. And crucially this warrant is absent in other kinds of inferences about agents' behavior. Given the fact that I have seen Jane have coffee every day at 7 a.m., I can safely infer that she will do the same tomorrow. What she aims to do with me plays no role in warranting this inference, however.

Crucially, this kind of instrumentally constrained inference grounds a distinctive kind of signaling behavior in humans. Sometimes, we perform actions such that, given the fact that we are committed to coordinating towards a shared goal, these actions must be interpreted as communicating about the world. If I am looking for the shovel, your pointing towards the shed communicates that the shovel is in the shed. Interpreted otherwise, your pointing would fail to play an effective role as a cooperative act aimed at advancing our goals. Just as I can infer that by going to the shed you are going for the shovel given your commitment to plant a tree with me, this same commitment warrants my inference that by pointing at the shed you are communicating that the shovel is in the shed. By the same token, uttering the words “the shed” may serve the same role when embedded in our cooperative engagement. And more radically, we will see, this kind of reasoning underlies the literal interpretation of the utterance “the shovel is in the shed.” The fact that you are committed to advancing our shared goals

warrants my interpretation that you are using this sentence literally—as opposed to sarcastically or in pretense.

In cooperative signaling, the fact that interlocutors are committed to advance a shared goal is a crucial premise in the speaker’s practical reasoning in issuing a signal to communicate a content and in her audience’s epistemic reasoning in interpreting said signal as communicating said content. Knowing that you know that I am committed to advancing our shared goals, I issue a signal that must be interpreted as communicating that the shovel must be behind the shed if it is to make sense as a move in our cooperative engagement. That same mutual knowledge gives you reason to interpret my actions as I intended. Ultimately, the distinctive rationality of cooperative action enables us to issue signals that are contextualized as providing information to advance the shared goals of the conversation.

This kind of signal distinguishes human from nonhuman communication. The pheromonal trace ants leave to each other carries information about the world that is already encoded in their genetic makeup—“this trail leads to food/danger/the queen.” These signals are more similar to how flies can tell that food is around because of the residual trace left by rotten fruit than to the ways humans can imbue utterances with meaning in conversation. Human cooperative signals make use of our mindreading abilities to decipher each other’s communicative intentions, given what we are doing with each other. This allows for an incredibly rich variation in conventional and non-conventional strategies to instance the cooperative signaling described above.

This, I submit, is the distinctive core of the speech act in question. To assert is to issue a cooperative signal that must be interpreted as aimed at communicating information about the world given the purposes of the conversation. Assertion—and its manifold of normative dimensions—is enabled and distinguished by how our interpersonal lives and abilities enable us to act with each other in complex and creative ways. In turn, since the act of assertion can only happen between agents that recognize each other as willing and able to act in these ways, whatever is required from an asserting speaker will inescapably involve other people.

Of course, there are a good number of evident complexities. For instance, what is the role of language in all this? Are diary entries, soliloquies, one-off statements excluded from the realm of assertion? How about road signs or a message in a bottle, destined for whomever happens to see them? Can cooperative signals even ensue when the speaker aims to deceive their audience or in contentious conversations like negotiations or conversations with hostile interlocutors? Hopefully, throughout the chapters I provide at least partial answers to these natural questions.

The Chapters

Here I give a bird's-eye view of the arguments of each chapter of the dissertation. Before, however, two cautionary notes.

First, the reader will find the unity of my overall argument somewhat repetitive. In each chapter, I aim to address a distinct philosophical problem by appealing to the act of cooperatively offering information as means to achieve a shared goal. Since they are meant to be read as standalone essays, each chapter has more or less the same section characterizing this act in terms of cooperative action.

Second, in another sense, there is a glaring disunity in the dissertation. Although I have been talking of assertion in this introduction, I only use the term 'assertion' in chapters 4 and 5. In the rest of the chapters, I use different terms (testimony, committal speech, addition to the common ground). Again, this owes to the fact that each chapter is a standalone piece engaging with a different literature and thus makes use of different terms. As a whole, however, the dissertation can be seen as making the case for coalescing all these terms. I say more about this in the next section. For now, I present the chapters.

Together, Chapters 1 and 2 can be seen as arguing for the perhaps unexpected thesis that language is not all that important to characterize the act of assertion. There is a widespread philosophical impulse to argue or simply assume that there are distinctive epistemic and moral norms reserved for the strict—semantically determined—use of language to transmit our thoughts. This literalist impulse, I argue, is misguided. There are pervasive and systematic ways of communicating nonliterally that are subject to the same epistemic and moral norms that govern literal speech.

Chapter 1 argues that there are no good epistemic reasons to think that testimony is constrained by literal communication. Chapter 2 argues that there are no good moral reasons to trace the distinction between committal and noncommittal communication along the distinction between literal and nonliteral communication. Rather, the relevant epistemic and moral distinctions should be understood as the distinction between contents that are required and contents that are merely allowed as cooperative moves in conversation.

Chapter 3 turns to the interface between the epistemic and interpersonal dimensions of testimony. When things go right, our testimony justifies other people's beliefs. When things go wrong, we entitle people to reproach us for offering defective testimony. These epistemic and interpersonal features of testimony are, respectively, the center of attention of evidentialist and assurance accounts in the literature. Neither of these families, I argue, can provide a successful explanation of why the epistemic and interpersonal features of testimony *come*

together. I argue that understanding testimony as a cooperative signal yields an account of this dual nature.

Chapters 4 and 5 turn to debates on what has been dubbed the norm of assertion. A venerable constitutivist tradition in epistemology aims to provide an account of assertion in terms of the distinctive norm that governs it. Assertion is the unique speech act governed by the norm of assertion. I think this approach is right. However, one must be careful in picking what kind of norm is at stake.

Chapter 4 argues that recent views that appeal to etiological functional norms are misguided, as they cannot provide a single norm that constitutively individuates assertion. However, functionalism is not lost if one appeals to the artifactual function of assertion, as it is given by the cooperative aim that distinguishes cooperative signals as described above. In chapter 5 I generalize this argument and hold that the norm of assertion cannot be game-like, conventional, social, or functional. Rather, the norm of assertion must be an instrumental norm, stemming from the rationality of cooperative action. In turn, I show how this approach yields a distinctive path to derive the content of the norm of assertion. It turns out that, given the doxastic requirements of intentional cooperative action, speakers should assert true, believed, justified, and safe contents. However, whether this is knowledge will depend on your epistemological allegiances.

What is Left Unsaid

I conclude this introduction by briefly addressing some ideas that, although present across the dissertation, are not explicitly defended in any of the chapters. As a heads up, this is the most speculative part of the dissertation. But it also marks some paths ahead for cooperativism as a speech act theory.

Assertion, Semantics, and Normativity. If I am right, there is good reason to think that there is a unified speech-act-kind grounded on the rationality of cooperation. And again, if I am right, this same act has received many names in philosophy: testimony, assertion, statement, assurance, addition to the conversational score, etc.

Correspondingly, throughout the dissertation, I use different labels to refer to this act—mainly ‘testimony’ (chs. 1, 3), ‘assertion’ (chs. 4, 5), and ‘committal speech’ (ch. 2). This unfortunate lack of terminological unity is due to the fact that each chapter of the dissertation is meant to be a standalone piece that engages with a distinct debate and thus with a different literature. These different pockets of philosophy use different terminology and, as I have learned, imposing my own terminology leads to confusion.

Hopefully, the entire dissertation will amount to an argument that we can coalesce our terminology and choose one label. In this introduction, I have been using the term ‘assertion’, as I prefer this label to refer to the act I am characterizing. In turn, I would suggest that ‘testimony’ refers to its epistemic aspect, ‘committal speech’ captures its moral dimension, and ‘addition to the conversational score’ is a technical term to model the dynamics of communication.

However, this choice reveals a tension between two uses of the term ‘assertion’. On the one hand, there is the idea that assertion is semantically constrained, and thus that asserted contents are closely constrained by the minimal proposition expressed by a contextually saturated declarative sentence. This is, for instance, the use some linguists give to the label, as it differentiates semantically determined content from other discursive levels, such as implicated, presupposed, or not-at-issue contents. On the other hand, philosophers and in particular epistemologists tend to use assertion to refer to a normatively significant speech act—a distinctive declarative speech act that is characteristically governed by epistemic and non-epistemic norms. A similar idea is endorsed by philosophers that have made assertion an important aspect of the ethics of speech—as something that can be subject to testimonial injustices or systematically silenced, for instance.

There are a good number of theorists who assume that these two concepts align. Thus, they assume that non-semantically constrained contents are just not subject to the norms of assertion. But in most cases, this is just that: an assumption. And if the arguments of chapter 1 and chapter 2 come through, we may do better by eschewing this assumption. There are no good epistemic or moral reasons to think that literal speech is normatively special. Moreover, if the arguments from chapters 3, 4, and 5 are correct, we have reason to think that the unified act of cooperatively communicating a content is subject to the epistemic and interpersonal normativity that is taken to characterize assertion.

There is much we can learn from investigation in semantics. Yet, we must be careful to not overplay semantics—to not fetichize language—when addressing normative questions without good argument for it. So, my suggestion is that we should keep these two concepts—semantically constrained contents and normatively significant declarative speech acts—separate. An important consequence of this is that some nonliteral contents, like fully cooperative implicatures, will be subject to the norms of assertion.

Of course, as a philosopher, I prefer to reserve ‘assertion’ for the normatively significant speech act of communicating information about the world, and live with the consequences of endorsing this revisionist, non-semantically-constrained notion of assertion. Not least because

there are perfectly good alternative terms to refer to the semantically determined content expressed by our sentences—e.g., semantic content, minimal content, or literal content. Yet, I struggle to find an alternative term for the normatively significant speech act that I am characterizing here.

Metaphysics. There are at least two ways of asking the question ‘What is the content of an assertion?’. For starters, and most obviously, one may be asking what content is expressed by a token assertion. That is, one may ask what message or content is being communicated by a token assertion. Yet, one may also be asking about what in the world amounts to an asserted content. Other ways to ask this second question include what kind of thing are asserted contents or what in the world are theories of assertion trying to model. Let me say something about cooperativism’s tentative answer to this latter question.

As a speech act theory, cooperativism lies between conventionalism and mentalism. The content of a token assertion does not reduce to the content of the semantic features of our uttered sentences. We often assert (in the normatively significant sense described above) contents beyond the conventional features of our uttered sentences—as in the use of fully committal implicatures (see chapters 1 and 2). However, the content of our assertions does not reduce to the mental states of conversing interlocutors either. Sometimes, for instance, we issue accidental assertions and communicate something we do not mean.

Asserted contents are public, objective, and normative at the same time. However, they are not conventions, intentions, or beliefs. How is this possible? Asserted contents are individuated by their role in explaining agents’ actions within a coordination engagement. The content of an assertion is the information that rationalizes an agent’s utterance as a cooperative signal—as an act aimed at providing information to achieve a shared goal. Why did Ana issue that utterance? Because she wanted to communicate that the sweater is in the suitcase.

There is precedent for this kind of action-first approach to meaning—one that goes back at least to Pierce and Ramsey’s pragmatism, extends to Putnam’s externalism, Dretske’s and Millikan’s functionalism, and is present even in Stalnaker’s conversational dynamics. In one way or another, these are views that individuate meaning in terms of what it takes to successfully act in the world.

While usually these are views that aim to individuate mental contents, my suggestion is that, as informational states, asserted contents are individuated by their role in explaining *joint* action. The content of an assertion is given by the relationship of the interlocutor’s shared goals and beliefs with a speaker’s utterance, such that said utterance is rationalized as aimed

to provide information to advance said goals. Asserted contents amount to the information that is rationalized as aimed to be given by the speaker given what she is trying to do *with* her audience. Once we see joint actions as a distinctive kind of action, we can see asserted contents as a distinctive kind of content.

So understood, asserted contents must be public, as they depend on how the interlocutors mutually acknowledge the world is and should be. No private mental state has a determining role in what is asserted, since no private content can shape mutual beliefs or aims. Asserted contents are also objective, for they are determined by the instrumentally rational relationship between an utterance, the shared goals of the interlocutors, and what the world demands from them to actually achieve these goals. And asserted contents are normative. In the good case, asserted information aligns with what the speaker intends and with how the world is. Yet, asserted contents can fail to fulfill this role. They may be misaligned with the speaker's intentions, the world itself, or both.

Crucially, this cooperativist view of asserted contents draws a straightforward line between content individuation and interpersonal normativity. Since asserted contents are enabled by the things we do with others, failed assertions are interpersonal failures to do what one committed to do with other agents. Our interpersonal lives are inescapable when accounting for the metaphysics of assertion.

Beyond Philosophy. The model I offer here has the advantage of playing well with other social sciences. It dovetails, for instance, with questions in psychology concerning our cognitive abilities to read each other's minds and act together in complex ways. Much work in evolutionary and social psychology, for instance, is concerned with the phylogenetic and ontogenetic stories of our distinctively human interpersonal skills. It should not be surprising that the distinctive ways we communicate with each other are deeply related to the distinctive ways we think about and with each other.

Further, insofar as joint action can be modeled by the rational interplay between interlocutors' private and public beliefs and preferences, the cooperativist picture is amicable to game theoretic formal methods. In fact, the cooperativist picture recommends formal methods that model the interaction between agents' individual and mutual beliefs *and* the intra- and extra- discursive goals that guide and motivate their conversations. Communicated contents should be modeled in terms of the role they play in non-language-specific interactions between interlocutors. After all, communication and language are seldom performed for their own sake; they imbue and characterize the many ways we act in the world.

While there are important formal tools in the literature that would allow for this kind of model, such a discourse dynamics model is still to be seen.

The cooperativist picture also dovetails with formal and non-formal models in sociology and anthropology concerning how social norms, institutions, and power shift our preferences and expectations in ways that constrain, mold, and motivate the reasons we act together and thus communicate with each other. Communities often generate higher-order social phenomena to ensure coordination and mutual action among their members. There are, for instance, complex systems of power and reputation that motivate agents to modulate their communal actions in ways that instance and maintain social hierarchies. Similarly, communities put in place systems of informal social sanctions (or social norms) that motivate agents to coordinate when they otherwise would not. And we often build institutions that mold and define new ways in which we act together.

If assertion and its plurality of normative dimensions are determined by how and why we act together and these higher-order social phenomena constrain, motivate, and mold the ways in which we act together, higher-order social phenomena must have a central role in a full theory of assertion. One such role will be to account how, in being a deeply social act, assertion can be subject to particular kinds of social injustice. Biases, prejudices, norms, and institutions can place systematic barriers that exclude agents from engaging in fully cooperative actions and thus from imbuing their assertions with the purposes they want. However, it is also the case that these social phenomena often motivate and force these agents *into* second-personal cooperative engagements that they would not endorse otherwise. Unjust threats from power and social sanctions, for instance, may coerce people into cooperative actions that motivate truthful assertions. In this latter sense, unjust social phenomena can hijack the instrumental normativity of assertion to preserve oppressive social conditions.

Ultimately, the picture I favor invites us to abandon an understanding of human communication as something a speaker does *to* an audience; to assert is not to encode a message sent to and deciphered by another agent. Rather communication is something we do *with* each other. We are able to communicate in virtue of the fact that we are the kind of beings able to distinctively recognize each other as willing and able to do things that would not be possible when acting alone. Assertion is fundamentally a kind of action that inescapably involves the mutual recognition of others as agents like us. No wonder, then, that it lies at the center of our moral, epistemic, and social lives.

CHAPTER 1

Communicating Testimonial Commitment

Abstract: I argue for the Cooperative Warrant Thesis (CWT), according to which the determinants of testimonial contents in communication are given by the practical requirements of cooperative action. This thesis distances itself from conventionalist views, according to which testimony must be strictly bounded by conventions of speech. CWT proves explanatorily better than conventionalism on several accounts. It offers a principled and accurate criterion to distinguish between testimonial and non-testimonial communication. In being goal-sensitive, this criterion captures the role of weak and robust cooperation in determining the contents to which speakers testify or fail to testify. And, finally, it yields a principled explanation of why testimony entails the epistemic commitments that distinguish it as an epistemic source.

1. Overview

Testimony, understood in the epistemologist's sense, is a ubiquitous epistemic source. Much of what we know (that Mars has two moons, that the bar on Green Street has stout on tap, that I have two kidneys) we know from accepting the testimony of others. Yet, although there is a plethora of ways to transmit information through speech, not all of them count as testimony. A superintendent who utters "I gave him a hundred thousand reasons to retire" may somehow communicate (suggest, insinuate) to his audience that he bribed the police commissioner. Intuitively, however, he has not testified to bribing anyone.

Explaining why the superintendent's utterance does not amount to testimony concerns the often neglected *communicative dimension* of testimony. To transmit testimonial knowledge about p , speakers must not only be in a privileged epistemic position regarding p , but they must also convey p appropriately. Independently of what our superintendent knows about the bribe, his utterance does not comply with the communicative conditions that fix testimonial contents. One may ask, then, what are the conditions that distinguish testimonial from non-testimonial speech.

My main contention is that the communicative dimension of testimony is best captured by appealing to the non-language-specific dynamics of cooperative action. This goes against the popular idea that that testimony must be strictly constrained by linguistic conventions. Our superintendent, the conventionalist story goes, did not testify because his communicated content goes beyond the conventional content of his uttered sentence. This view, however, is misguided. It fails to capture genuine testimonial communication that goes beyond the conventional features of our utterances.

We do better, I argue, by modelling testimony as the act of cooperatively offering information. Given a speaker's cooperative commitments, her audience is warranted in

interpreting some of her actions as serving the cooperative role of informing about the truth of a proposition (i.e., as testimonial utterances.) Crucially, this role can be fulfilled by conventional and non-conventional communicative strategies. Ultimately, then, our superintendent does not testify to his bribery because his utterance does not fulfill the appropriate cooperative role in conversation. Call this view the Cooperative Warrant Thesis (CWT).

CWT yields a better picture of the epistemic grounds that warrant the interpretation of testimonial contents. First, it systematically captures a more accurate range of communicative strategies that can fulfill a testimonial function while excluding non-testimonial speech. Second, being goal sensitive, this view captures the modulating effects of weak and robust cooperation on testimonial contents. What communicative strategies count as testimony can differ between a conversation with an old-time friend and a conversation with an adversarial crime suspect. This modulation effect, so far overlooked in the literature on testimony, accommodates the initial appeal of conventionalist intuitions. And third, as opposed to conventionalism, CWT provides a principled explanation of why our testimonial practices involve epistemic commitments.

In the next section, I argue that testimony can be pragmatically determined in ways that conventionalism cannot trace. In §3, I characterize a distinctive kind of publicly available interpretive warrant offered by cooperative action in general. In §4, I explain how this cooperative warrant characterizes testimonial contents and how it applies to cases where testimony does not seem to involve cooperation. In §5, I turn to explain why some kinds of declarative communication do not count as testimony and how the threshold between testimonial and non-testimonial communication can shift with a conversation's degree of cooperation. In §6 I offer a principled explanation of why the practical normativity of cooperation, and not convention, grounds the epistemic commitments of testimony.

2. The Dynamics of Testimony

Testimony is a second-personal epistemic practice. Unlike speculating or guessing speakers, testifying speakers epistemically commit in a stronger, distinctive way to a communicated content.¹ When a speaker testifies to her audience, she commits to having the appropriate epistemic privilege with respect to the content of her testimony.² Among other features, the epistemic commitment that characterizes testimony entitles audiences to hold testifying

¹ Although there are crucial differences in the role this fact plays in a full account of testimony, it is widely accepted that testimony constitutively involves practices of interpersonal commitment. See, for instance, Adler (2002), Coady (2000), Cohen (1982), Fricker (1994; 2006; 2012), Goldberg (2007; 2012), McMyler (2011), Moran (2018; 2013).

² What kind of epistemic privilege testifying speakers must have has been widely debated in the literature—for instance, certainty (Stanley 2008), knowledge (Williamson 1996), justified belief (Kvanvig 2009; Lackey 2008; McKinnon 2015), or rational belief (Douven 2006). I remain neutral about this aspect of the epistemic dimension of testimony.

speakers accountable for the epistemic merits of what they say. When speakers offer defective testimony, audiences are at the very least entitled to ask speakers for the reasons they had to testify and reproach them if these reasons are not excusing.

For our purposes, then, the relevant explanandum is that some means of communication carry the epistemic commitments that characterize testimony and others do not.³ Consider our superintendent. It would seem that he has not testified that he bribed the police commissioner precisely because he cannot be pinned down for having committed to that content. No matter what he knows about his bribe, he seems able to validly deflect claims of accountability for having communicated that he bribed the commissioner (“I never told you I bribed him!”). In contrast, the interpretation of testimonial contents must be warranted in a way that blocks this kind of deflection. Hence, an account of the communicative dimension of testimony should explain why testimonial communication involves interpretive warrants that epistemically secure communicated contents in a way the superintendent’s speech act does not.

At first glance, it may seem natural to appeal to the fact that what the superintendent communicates simply goes beyond the semantic features of his utterance. Hence, some may suggest, testimonial contents must be semantically constrained. Elizabeth Fricker (2012) and Sanford Goldberg (2007; 2008) have independently defended such a view.⁴ Testimonial contents, they argue, strictly trace the syntactic-semantic features of our uttered sentences and allow only for minimal supplementation of context-sensitive expressions needed to arrive at a truth-evaluable proposition. Their views allow, for instance, fixing the referent of pronouns like ‘she’ to refer to Jane Smith and disambiguating terms like ‘bank’ to refer to the financial institution. Nevertheless, any further enrichments or alterations—any “logical, nomological, or conversational implication” (Fricker 2012: 97)—must be excluded from the realm of testimony. For ease of exposition, let us call the resulting contents the literal contents of our words.⁵

³ My focus on commitment differentiates the question I am asking from Peet’s (2019) question on what it takes for communication to be knowledge-yielding. While Peet is concerned with what it takes for communication to yield knowledge, I am asking what it takes for a speech act to count as testimony and epistemically commit speakers to a content—independently of whether knowledge has been transmitted. For all we have said, there can be non-testimonial, knowledge-yielding communication, and testimonial non-knowledge-yielding communication.

⁴ Although without developing an argument for it, other philosophers explicitly endorse this view, see Burge (2013: 248, 355–56) and Owens (2017: 215). Other philosophers seem to implicitly endorse it (MacFarlane 2011).

⁵ Literal contents can be thought of as close to what has been called minimal contents (Borg 2010; 2019; Cappelen & Lepore 2005). Nonetheless, whether there are such contents is famously a contentious issue. See Bach (1997; 2005), Carston (1999), Korta and Perry (2006), Recanati (2003), Stanley and Szabó (2000), Stokke and Schoubye (2016).

There is *prima facie* good reason to accept this view. Under a widespread construal, non-literal enrichments are determined by speaker intention.⁶ However, agents have special epistemic authority over their own mental states. Arguably, then, agents are able to validly deny having a given intention without much space for others to show evidence that they indeed had that intention. In contrast, literal contents are determined by public conventions of speech shared by a linguistic community and can be thus appealed by audiences in holding speakers accountable. Insofar we are interested in characterizing an epistemic practice that grounds epistemic commitments against which speakers can be held accountable, semantic convention seems to be a safe choice.

However, literalism about testimony is excessively restrictive. Pragmatic enrichment is just too pervasive in our everyday communicative practices. Picture a familiar scenario:

Charcoal: María and Gabriel are planning a barbeque. It is clear that they have everything they need but charcoal. Both know there is only time for one trip to the store. Gabriel offers to go, but he is from out of town, so María says:

(1) There is a store down the road.

Although it goes beyond the literal content of (1), María is arguably committing herself to there being an open store down the road where Gabriel can buy charcoal. As an initial indication, notice how the following two possible responses from Gabriel would be redundant (as indicated by the hashtag):

(2) #Is it open?

(3) #Do they sell charcoal?

These questions are redundant insofar as they ignore the role (1) has in conversation. If María's utterance is to make sense as an input in the conversation is because the store is open, and it sells charcoal. To use a widely accepted framework, the fact that these questions are redundant strongly suggests that this information has entered the publicly available conversational record (cf. Stalnaker 2014; Roberts 2012; 2018).

Crucially, there is no need to appeal to María's private intentions to discern her commitment. That the store is open and it sells charcoal follows from María and Gabriel's common context, the fact Gabriel offered to go to the store, and, crucially, their avowed common interest in acquiring charcoal for their barbeque. If it turns out that the store is closed or that they do not sell charcoal, Gabriel can appeal to these public features of the conversation

⁶ See, for instance, Schiffer (2006), Bach and Harnish (1979), Neale (2004), Harris (2019). Arguably, these authors have their own answers to anti-intentionalist complaints. Below I propose a solution that focuses on the epistemic warrants provided by cooperative reasoning, rather than speaker intention. See n. 16.

to hold María accountable. She cannot plausibly deny they were planning a barbeque or that Gabriel offered to go to the store.

Admittedly, it may be that María genuinely thought the store was open. What matters, however, is that her appropriate response to Gabriel's reproach should be epistemic, not communicative.⁷ She should excuse her commitment to inaccurate information ("I thought it was open"), rather than deny having committed to such information ("I never said it was open").⁸ Further, if there is indeed charcoal at the store, Gabriel forms the corresponding belief, and María is competent and reliable, Gabriel can gain knowledge from her utterance. If so, it is not clear why it should not count as testimonial knowledge.

It would seem, then, that María's commitment to the non-literal content of her utterance is determined by something beyond the literal content of her words and also independent from her mere private intention. Further, her commitment to her communicated content is fixed in a way that our super intendent's insinuation is not. In the next sections I present a view that make sense of these features. First, however, let me consider a possible conventionalist retort.

Semantic-syntactic conventions do not exhaust linguistic conventions. Many theorists model conventional meaning beyond the strict conditions Fricker and Goldberg have in mind.⁹ Grice, for instance, held that the causative reading of 'and' in utterances like "The Dow fell and investors panicked" owed to a convention beyond the literal meaning of these words. So, those wary of the restrictiveness of literalism but convinced by the power of convention may argue that testimony is bounded by a broader set of discursive conventions.

This move, however, will not go far. Shifts in context can modify testimonial commitments in ways that cannot be traced even broader discourse conventions. Picture a variant of the above scenario.

Charcoal*: Jane and Sam are planning a party. However, they disagree on what kind of party to host. While Jane wants to host a barbeque, Sam wants to host indoors. It is public to both that while either party requires food and drinks, a barbeque would

⁷ The specifics of the conversation may make it the case that the appropriate warrant to testify in this case is less strict than in other cases (cf. Goldberg 2015). Even so, the case suggests that María would be accountable for having the same warrant had she uttered, "there is an open store down the road, and it sells charcoal."

⁸ This is an important result. At least in this case, the fact that she can much later point out that the content she committed to was not semantically or logically entailed by the literal content of her words does not dissolve her epistemic commitments. Mere cancellability does not amount to plausible deniability. See Welker (1996) and Roberts (2012) for similar arguments.

⁹ See, for instance, Lepore and Stone (2015a; 2015b), Abusch (2010), Van Der Sandt (1992), Asher and Lascarides (2005), Hobbs (1985).

require charcoal, and an indoor party would require decorations. There is only time for one trip to the store and Sam offers to go. Sam is from out of town, so Jane says:

(4) There is a store down the road.

Notice how the commitments at stake shift. Take our test questions:

(5) #Is it open?

(6) Do they sell charcoal?

Given the scene, (5) is still redundant. Jane and Sam are committed to hosting a party, so (4) does commit Jane to the store being open. Yet, (6) is no longer redundant. Given Sam and Jane's disagreement on what kind of gathering to host, Sam is no longer entitled to interpret (4) as committing Jane to the fact that the store sells charcoal; even when both know that charcoal is required for a barbeque. That (6) is entirely appropriate, hence, suggests that the store selling charcoal has *not* publicly entered the conversational score and, therefore, that (4) does not carry the same commitments (1) did.

Conventions do not seem to be the relevant content-determining factor here. What determines the content to which Jane testifies is best explained by the instance-specific, pragmatic interaction of her utterance with the goals of the conversation and the information public to her and Sam. Lenient conventionalists may be tempted to suggest that the content of (4) is somehow guided by a conventional discursive (though non-semantic) relation that fixes the content to which Jane commits.¹⁰ This suggestion, however, is unwarranted. To prove useful, it must flesh out how the specific commitments entailed by Jane's utterance precisely trace practices accepted and deployed by her and Sam's linguistic community.¹¹ Until that suggestion is developed, the onus is still on the conventionalist.

Although the idea that conventions and commitment share a deep conceptual connection has been defended even outside the epistemology of testimony (Camp 2019; Kölbel 2011; Stainton 2018), it seems worthwhile exploring alternative answers. I now turn to such an account. In §6, I come back to the conceptual question and argue that appealing to the practical normativity of cooperative action, as opposed to convention, offers a better explanation of why testimony and commitment go together.

3. Cooperative Warrant

As our superintendent example shows, not every kind of communication yields testimonial commitments. Fricker and Goldberg are right in suggesting that the determinants of

¹⁰ Lepore and Stone (2015b: 87) suggest this, citing Asher and Lascarides (2005). Nevertheless, they do not offer much detail about this proposal.

¹¹ See Simons (2018) for arguments against this suggestion.

testimonial commitments must ensure safe and publicly warranted interpretations that can be appealed to when holding speakers accountable. We have seen, however, that conventions alone cannot account for this. We do better, I submit, by delving into the practical structure of non-language-specific cooperative action.

In a sense, this idea is not new. It is the key insight of Grice's (1967) Cooperative Principle, and lies at the core of a wide slew of theories of communication. What is distinctive of what I have to say is the insistence that cooperative reasoning itself offers the adequate *epistemic* grounds to secure testimonial contents. Crucially, at least when it comes to testimony, this means that we should abandon the classic idea that non-literal contents are directly determined by speaker intention. In this section I explain how cooperation in general offers a distinctive kind of publicly available warrant for action interpretation. In the next two sections, I turn to explain how this warrant accounts for the distinction between testimonial and non-testimonial communication.

Two (or more) agents have a *common goal* when each is individually committed to actualizing the same world-state. In turn, two agents have a *shared goal* when they are mutually committed to actualizing the same world-state. Adapting Bratman's (1992) classic example, all it takes for you and me to have the common goal of painting a house is that we are committed to having it painted. We share that goal when we commit to doing so together—when we commit to coordinate and interlock our actions in a mutually responsive way that leads to having the house painted.

Like individual goals, shared goals offer *inferential warrants*. Just as the individual intention to ϕ allows me to take for granted 'I will ϕ ' in deliberation, cooperating agents regularly appeal to the fact that they are committed to a shared goal in their reasoning about their and their counterpart's actions.¹² In other words, the fact that agents strive towards a shared goal can be used as a premise to infer the role of their counterpart's actions in achieving that goal. If you reach for the roller instead of the brush, the fact that we share the goal of painting the house warrants my belief that you will paint the house's walls.

Unlike individual goals, shared goals involve *publicity requirements*. Agents who commit to a shared goal commit (on pains of procedural failure) to act in ways that sufficiently disclose the roles of their actions as means to achieve their shared goal.¹³ Suppose we have jointly committed to painting the house and I stop and leave. I must make it clear that I am going for more paint to avoid you inferring I am no longer jointly committed to painting the house and this dissolving our shared goal—even if I indeed had the intention to go for more paint.

¹² This has been defended in Bratman (1993; 2014), Cohen and Levesque (1988; 1991), Tuomela (2006).

¹³ Compare with Bratman (1993; 2014), Cohen and Perrault (1979), Gilbert (2009)

Together, publicity requirements and inferential warrants give rise to a distinct pattern of inference about action. In cooperation, beliefs can be substantiated by inferences that rationalize an action under the assumptions that cooperating agents are (i) in fact engaged in a shared goal; and (ii) committed to disclosing the role their actions have as means towards said goal. Call the propositions that serve this explanatory role *cooperatively warranted propositions*.¹⁴

Consider an example before I offer a more detailed definition. Say that you and I have committed to painting the house, but I stop and go to a nearby shed. Say too that it is public to both of us that there is where we store the paint. Our commitment to cooperate yields a particular pattern of reasons to believe that I went to get more paint. For starters, this conclusion is warranted given our shared goal and the fact that it is public to us that there is more paint in the shed. Further, it is also warranted by the fact that I am committed to disclosing the role my actions have in achieving our goal. It is public to both of us that going for more paint is the best explanation of my action. Hence, were I not going to get more paint, I would have disclosed why else I am going there. The fact that I did not do so gives you further reason to believe that I am going for more paint. The proposition that I went to get more paint is therefore cooperatively warranted.¹⁵ The interaction between my commitment to our shared goal and my action makes it epistemically permissible for you to treat the proposition at stake as true.

We can now be more schematic. Given an agent X, an action ϕ , a proposition p, and an acknowledged shared goal SG,

Cooperative Warrant: When the truth of p is required by the best explanation(s) of how X's ϕ -ing is a means towards SG, given information public to the agents engaged in SG, p is cooperatively warranted.

In the sense relevant here, the best explanations of an action provide the most stable, coherent, simple, and reliable account of how that action is a means to a shared end.¹⁶

Since cooperative warrants rely on the assumption that agents are committed to disclosing the role of their actions, the explanations that yield cooperatively warranted contents must be grounded on public information. If only you or I (and not both) knew that there is extra paint in the shed, your belief that I went for more paint would not be cooperatively warranted—even if you did arrive at it through other means.

¹⁴ Note that this does not require positing ontologically distinct collective beliefs as in Gilbert and Priest (2013) to explain access to warranted beliefs about what is communicated.

¹⁵ I say that a proposition, instead of a belief, is cooperatively warranted to mark that our cooperating generates reasons to believe the proposition that I went for paint even if you did not come to believe so.

¹⁶ Cf. Harman (1990), Bratman (2014), Cohen and Levesque (1991).

It may well be that two or more explanations are equally fit to explain a cooperative action. If so, only the propositions required by the intersection of all the best explanations are cooperatively warranted. If we both know that there are brushes and paint in the shed, the proposition that I went for both or either of them (but not for one in particular) will be cooperatively warranted.

Importantly, cooperatively warranted propositions can differ from an agent's intention in action. Given my cooperative commitment to paint the house with you, the cooperative interpretation of taking the brush is that I aim to paint the house—even if for some reason I aim to paint my car. Of course, once I start painting the car, you can conclude that I am no longer committed to painting the house. That, however, does not mean that your original interpretation was not cooperatively warranted.

Lastly, doubts about cooperative commitments cast doubts on cooperatively warranted propositions. Your belief that I went for more paint will be put into doubt if you question whether we are in fact cooperating towards painting the house. By the same token, sometimes, a cooperatively warranted belief will be defeated by evidence that one's counterpart has defected from the shared goal. If I go to the shed and do not come back, the cooperatively warranted belief will be defeated in favor of the non-cooperative explanation that I no longer share the goal of painting the house. Nevertheless, whenever it is evident that our counterparts are committed to a shared goal with us, cooperative warrants serve a vital role in explaining each other's actions in ways that enable complex forms of coordination.

4. Testimonial Communication

I have argued that the practical pressures involved in cooperation warrant a distinctive pattern action interpretation. I turn now to argue that this kind of warrant grounds the communicative dimension of testimony.¹⁷

When two (or more) agents share a goal, some actions must be cooperatively interpreted as aimed to actively inform about the truth of a proposition.¹⁸ Put otherwise, to make sense of some our counterpart's actions as aimed towards a shared goal, it is often required to interpret them as aimed to offer information relevant for achieving said goal. My uttering "There is more in the shed" can be interpreted as aiming to inform you that there is more paint in the shed because I embed it in the cooperative scheme of painting the house. That the purpose of my utterance is to share that information is clear because of my

¹⁷ Note that the following view is in line with Greco's (2020) point that cooperation is fundamental to testimonial exchanges. Nevertheless, Greco does not address the communicative dimension of testimony and I do not endorse the epistemic consequences he sees stemming from this fact.

¹⁸ Agents can cooperatively communicate with others who they do not know, as when I tell someone behind a curtain that the theater is closed; potential cooperators in the future, as when I leave a note that the fridge is not working; and even groups of people, as when I announce that dinner is ready.

commitment to paint the house and my commitment to paint it *with* you. The reasons I have to act communicatively and the reasons you have to interpret my action as communicative hinge on the fact that we are jointly committed to a shared goal.

That said, we can understand certain signals as warranted by the cooperative features of a conversation. Given a proposition *p*, an utterance *U*, a speaker *S*, an audience *A*, and shared goal *SG* acknowledged by *S* and *A*:

Cooperatively Warranted Signal: When interpreting *U* as informing *p* is required by the best explanation(s) of how issuing *U* is a means towards *SG*, given information public to *S* and *A*, *U* cooperatively signals *p*.

The relevant cooperative warrants are provided by the fact that an agent's cooperative action is best interpreted as providing information about the world. In the sense at stake here, a proposition is cooperatively signaled by an utterance when it provides the most coherent, simple, and reliable explanation of a speaker's communicative action as embedded in a cooperative interaction.

My contention is that testimonial contents are cooperatively warranted signals. Thus, the following thesis:

Cooperative Warrant Thesis (CWT): Speakers testify to the propositions cooperatively signaled by their utterances.

Two clarifications are in order. First, CWT is meant to capture the epistemic grounds that secure testimonial interpretation. Hence, it is ecumenical with respect to how cooperative inferences should be modeled or whether cooperatively warranted signals amount to other discourse categories in the literature—like sayings, tellings, statements, or assertions.¹⁹

Second, the sense of cooperation at stake is permissive. All agents need to testify is present themselves as committed towards a shared goal and acting mutually and responsively to achieve it. This allows for testimony to occur in non-ideal cases. Speakers, for instance, can acknowledge shared goals in bad faith. Deception is the paradigmatic case. Nonetheless, if deception is to be effective, the features of the conversation must indeed cooperatively warrant

¹⁹Arguments as to how testimony is best modelled and whether it amounts to other categories in the literature extend the scope of this paper. Nevertheless, CWT is compatible, for instance, with Gricean maxim-centered models, questions under discussion frameworks, and game-theoretic pragmatics. Likewise, CWT makes no commitment towards what are the actual cognitive mechanisms that allow for content recovery. However, it has been shown that goal recognition and goal-based inference are crucial for linguistic interpretation in a way that is also orthogonal to conventional and non-conventional communication. Goal recognition is crucial for semantic supplementation (King 2014), semantic and discourse disambiguation (Hanna & Tanenhaus 2005; Hobbs et al. 1993), and diverse kinds of pragmatic enrichment (Roberts 2018; Simons 2010; 2018; Thomason 2003).

the testimonial role of an utterance.²⁰ Successful deceivers must offer good (albeit ultimately devious) indication that they are cooperative speakers. Further, cooperation can be variably motivated. You may paint the house with me because of our friendship or out of fear of what I would do if you do not. Hence, speakers can be coerced into testimonial cooperation through effective threats.²¹ In the next section, I say more about how contentious and ambiguous contexts modulate and constrain what counts as testimony.

Crucially for our purposes, CWT is orthogonal to the distinction between conventional and non-conventional communicative strategies. The fact that a string of words is conventionally associated to a content is not enough for it to serve a testimonial function in a given conversation. It has to be embedded in the right context that warrants its interpreting it as offering testimony. This is what distinguishes literal communication from irony, sarcasm and, more radically, cases of sleeptalking. CWT holds that the right context will be determined by the shared goals and public information that guide a conversation. Under these conditions, sometimes conventional contents will be cooperatively warranted and thus serve a testimonial function. If María utters “The store down the road is open and it sells charcoal” she testifies to the literal interpretation of her words because it is cooperatively signaled by this utterance.²²

Other times, as long as their utterances serve as cooperatively warranted signals, speakers can appeal to non-conventional means to testify.²³ In **Charcoal**, interpreting María as communicating that the store is open and sells charcoal rationalizes her utterance as a cooperative move towards her acknowledged shared goal. The cooperative role of (1: “There’s a store down the road”) is narrowed down to providing information relevant to obtaining charcoal. Once Gabriel and María have committed to the goal in question, there being charcoal in the store follows from a goal-grounded cooperative interpretation of her utterance and thus her utterance fulfills a (non-literal) testimonial function.

²⁰ Likewise, speakers can accidentally issue cooperatively warranted contents (*pace* Owens 2006). In the same context, María’s accidental uttering of (1) will still communicate that the store sells charcoal even if for some reason she would have preferred not to communicate this.

²¹ Whether testimony so motivated is reliable or not in these cases is an independent issue. What matters here is whether the cooperative features of the conversation warrant interpreting such an utterance as testimony.

²² Searle complains about Gricean accounts by pointing out that, in some cases, speakers communicate literal interpretations without intending it or even knowing their meaning (Searle 1965; but see Schiffer 1972). Under CWT, if a literal interpretation counts as testimony at all, it is because, under the given context, that is the only cooperative interpretation available.

²³ Although I focus on implicatures, cooperative signals can be enacted by a wide set of communicative strategies. This includes, for instance, non-verbal signals—like pointing to the garage when asked where the brushes are. Although these signals will require more from context to ensure a cooperative warrant, nothing about CWT precludes them to serve a testimonial function.

Under CWT, testimonial contents have external determinants that do not hinge on the mental states of any sole interlocutor.²⁴ The cooperative interpretation of María's utterance does not hinge on premises about her private intentions but on information available to both her and Gabriel and their shared commitments. María's publicly acknowledged cooperative commitment, together with her publicly available utterance, yields the cooperatively warranted interpretation that the store is open, and it sells charcoal—even if for some reason she actually intended to communicate something else with her utterance.²⁵

This explains why Gabriel has the standing to hold María accountable. He can appeal to the public features of the conversation that warrant María's cooperative signal to hold her accountable if it amounts to defective testimony. Of course, that a content is cooperatively warranted does not mean that it must in fact be added to the conversational record. Once recognized as such, cooperatively warranted interpretations can be rejected or challenged by interpreting audiences. Gabriel may still refuse to accept that the store sells charcoal. Nevertheless, that a proposition is cooperatively warranted fixes it as the testified content at stake.

The dynamics of cooperation also explain the commitment shift in **Charcoal***. In this case, Jane and Sam are jointly committed to hosting a party, which requires food and drinks. Nonetheless their respective individual goals preclude them from being jointly committed to hosting a barbeque (which requires charcoal) or an indoor party (which requires decorations). At most, then, they can be jointly committed to hosting a party that requires food and drinks, but not one that requires charcoal or decorations. Hence, the cooperatively warranted interpretation of Jane's utterance (4: "There is a store down the road"), includes only what is relevant for this limited shared goal. That there is charcoal in the store is *not* cooperatively warranted, even if it is required for Sam's individual goal. Likewise, that there are decorations at the store is not cooperatively warranted, even if it is required by Jane's individual goals. Even if Sam infers that Jane believes the store sells Charcoal, this inference does not depend on their shared commitment and therefore is not cooperatively warranted by the interaction. This is why his question (6: "Do they sell charcoal?") is still appropriate. Jane thus testifies only to the fact that the store is open, and it sells food and drinks.

²⁴ Compared to other Gricean accounts of communication, the distinctive aspect of CWT is that, at least when it comes to testimony, meaning determination is given by the rational interplay of utterance and cooperative commitment, not by the intention with which the speaker issues a specific utterance. In the successful case, intention and cooperative interpretation align. Nevertheless, cooperative interpretations can still stand in the absence of the corresponding speaker intention. Hence, the intention in action with which a speaker emits an utterance is not the intention that determines her accountability. Rather, what enables testimonial accountability is the expressed intention to achieve a shared goal, against which utterances are to be interpreted.

²⁵ Note that this renders CWT immune to objections of Humtdumptyism (MacKay 1968; but see Donnellan 1968).

Before I move on, let me address a possible concern. Intuitively, some agents can gain testimonial knowledge without sharing the goals of the speaker. This is paradigmatically exemplified by eavesdroppers (Lackey 2008; Owens 2006). Say that Ryan listens to María's interaction with Gabriel and acquires knowledge that there is charcoal at the store. To make the case stronger, we can imagine that María does not want Ryan to know there is charcoal in the store. In that case, Ryan could arguably gain testimonial knowledge without sharing any goals with María. Perhaps, then, cooperation is not essential to testimony.

Here, it is important to note that for a content to be cooperatively warranted, CWT requires that a speaker and an audience share a goal that rationalizes the speaker's utterance. That a third person witnesses the conversation does not bear on whether this requirement is met. Arguably, if Ryan gains testimonial knowledge at all, it is only in virtue of the fact that he understands María is indeed testifying to Gabriel that there is charcoal in the store. This is in turn only possible if Ryan understands the goals at stake and how they yield this content. Under CWT, this means that he understands that the cooperatively warranted interpretation of María's words is that there is charcoal at the store.²⁶ Testimony ensues when interlocutors cooperate towards a shared goal, even if this does not involve every other agent that can learn from it.

Summing up, modeling testimonial contents as cooperatively warranted signals captures the characteristic features of testimonial communication while allowing for the commitment shifts we observed in §2. Cooperatively warranted signals hinge on conversational features that are public to both speaker and agent and, thus, can be appealed to when holding speakers accountable for their defective testimony. I turn now to explain how CWT excludes non-testimonial contents and capture the effects of weak and strong cooperation in testimonial exchanges.

5. Non-Testimonial Communication

If CWT is correct, testimonial contents are fixed by the cooperative interpretations of an utterance, given the interests that structure a conversation. What, then, distinguishes non-testimonial communication, as exemplified by our superintendent's utterance?

Communication theorists have classically focused on conversations in which interlocutors' goals fully overlap, and utterances explicitly advance them. Yet, not all

²⁶ This also explains how one may gain testimonial knowledge from diaries (Lackey 2008; Owens 2006). Sometimes, diaries can be understood as means for an agent to cooperate with her future time slices. In those cases, an utterance in a diary may count as testimonial. If I read Raul's diary and know he uses it as a log for his house chores, I can gain testimonial knowledge that there is a spider nest in his attic. I'm in the same position of the eavesdropper. Decontextualized diary entries do not count as testimony. This is not a costly bullet to bite, for I can still treat the diary as a source of evidence, just not testimonial evidence (Coady 1992).

conversations are so fortunate.²⁷ Conversation requires that interlocutors to acknowledge *some* shared communicative interests. All in all, agents with no shared goals have no reason to even begin a conversation. Nevertheless, the relationship between utterances and goals can be ambiguous in ways that constrain cooperatively warranted contents and, therefore, testimonial commitments.

The superintendent's audience may want to know precisely how he convinced the police commissioner to retire. Yet, the nature of the interaction offers them evidence that the superintendent does not have the goal to share self-incriminating information. At most, then, the superintendent's audience can be confident that they share the unspecified goal of issuing relevant information about the police commissioner's retirement.

Under this goal, however, the utterance

(7) I gave him a hundred thousand reasons to retire

yields a wide set of plausible cooperative interpretations. And, while this set includes that a bribe was involved, it also includes alternative and innocuous interpretations to which the superintendent can appeal in denying his commitment—like having given the police commissioner many non-criminal reasons to retire. Although allowed as possible cooperative interpretations, none of these interpretations is required by the best explanations of (7) as means towards the shared goal of informing relevant information about the commissioner's retirement. Consequently, none of them is cooperatively warranted.

Unlike the testified contents in **Charcoal** and **Charcoal***, given the vague shared goal that guides the conversation, (7) communicatively fails to commit the superintendent towards a sufficiently specific content and, therefore, does not yield testimony about precisely how he convinced the commissioner to retire. While testimonial communication is determined by how a given interpretation is required by the shared goals that structure a conversation, non-testimonial communication exploits a range of allowed (but not required) interpretations to generate cooperative ambiguity. This ambiguity generates the uncertainty that characterizes insinuated speech (see, for instance, Pinker 2007; Asher & Lascarides 2013; Camp 2019). Things would be different if our superintendent uttered (7) in a conversation with his boss, who had instructed him to bribe the commissioner. Arguably, in that case, (7) will cooperatively warrant the content that the commissioner was bribed, and his utterance amount to testifying so.²⁸

²⁷ Recent years have seen exceptions to this trend. Some examples are Asher and Lascarides (2013), Saul (2019), Camp (2019). See Beaver and Stanley (2018) for a review of this perspective change in the philosophy of language.

²⁸ Say that this conversation between the superintendent and his boss has been bugged by the police. Since most probably the police do not share the information that is public between the criminals, they do not have enough

Of course, non-testimonial insinuations may occur in less contentious interactions. Two gossiping friends may be mutually interested in sharing information about Gale's love life. One of them may insinuate he is seeing something by uttering "well, he is going to the gym a lot lately". In this case the utterance at stake may play some cooperative role in the conversation. Nonetheless, it still fails to play the role of communicating *the* cooperatively required content that Gale is dating someone. If an utterance is to play the deniability-preserving role of insinuation, it must be vague enough to generate uncertainty among possible interpretations, given the acknowledged goals of the conversation. Such ambiguity may come from the indeterminacy of the goals or the vagueness of the utterance.

This marks another theoretical perk of CWT. It not only captures intuitions about epistemic commitments in cases like **Charcoal** (and **Charcoal***) and our superintendent's, but also accommodates intuitions that motivate conventionalism. In some cases, conventional moves in conversation can yield testimonial commitments, while non-conventional moves cannot.

To see how, note that had the superintendent uttered

(8) I bribed the police commissioner

He would have confessed in a way that does commit him to his bribery. Even under an unspecified shared goal, the cooperatively warranted interpretation of (8) is that he bribed the police commissioner. Given the widely accepted conventional features of (8), there is no room for alternative interpretations. This literal interpretation is the best explanation of his utterance as a means towards the goal at stake.

More generally, there is an inferential tradeoff between goal definiteness and conventional explicitness. We have seen that evident and specific shared goals allow pragmatically rich contents to be cooperatively warranted and serve a testimonial function. The fact that María and Gabriel have specified their shared goals to obtaining charcoal for the barbeque allows (1) to play the cooperatively warranted role of communicating that the store is open and sells charcoal. Nevertheless, unspecified goals make pragmatically rich speech cooperatively inert. Had María and Gabriel shared the goal of hosting any kind of party, (1) would not have yielded the cooperatively warranted content that there is charcoal at the store.

In that case, however, María could rely on the more conventionally explicit utterance

(9) You can buy charcoal in the store down the road.

evidence to account for what is cooperatively warranted in their conversation. Hence, speakers may testify only to a subset of their audiences.

Even under an unspecified goal, the conventional features of (9) narrow the relevant explanations in a way that conveys what (1) conveyed under a more specific goal. Against unspecified or uncertain goals, conventionally explicit utterances can narrow down plausible interpretations so that the only cooperatively warranted content is their conventional interpretation.

It is to be expected, then, that in contexts where it is not clear whether interlocutors share sufficiently similar goals (like police interrogations or contentious negotiations) cooperatively warranted contents will often converge with the literal interpretations of our utterances. Contexts like these tend to provide only evidence for unspecified shared goals. Therefore, widely shared linguistic conventions will tend to offer the safe, cooperatively warranted interpretations that yield testimony. This is why, in uncertain contexts, it makes sense for audiences to pressure speakers into using more explicitly literal utterances that exclude a broader set of possible interpretations.

Nevertheless, this effect is explained by the goal dynamics of the conversation, not by the mere historical record of conventions. As we have seen, in cases where there is robust evidence of sufficiently specific goal overlap, speakers can thoroughly and genuinely commit to contents yielded beyond the conventional meaning of a speaker's utterance. There is no need to exclude such cases from the realm of testimony just because in some other cases the cooperative features of a given conversation do not warrant non-conventional contents. This kind of modulation effect of cooperation on testimonial contents is neatly explained by CWT and is simply missed by conventionalism.

Moreover, CWT captures the fact that in some contexts not even Fricker and Goldberg's literal contents can fulfill a testimonial function. Insofar as literal contents require some contextual supplementation, interactions where goal overlap is extremely minimal will entitle audiences to reject even literal interpretations. These cases, for instance, may require making explicit demonstrative reference ("who do you mean by 'he'?"), relative adjectives ("how *tall* do you mean the suspect was?"), relational expressions ("what are you ready *for*?"); and disambiguation ("do you mean the suspect is in Geneva, New York, or Geneva, Illinois?").

Finally, CWT accommodates seemingly incompatible experimental data. Mazzarella et al. (2018) and Sternau et al. (2015) have found that what I have called here 'literal' and 'non-literal' conversational moves can enact communicative commitments for which speakers are held accountable. Yet, both studies also found that, comparatively, audiences tend to judge non-literal contents as less committing. However, the paradigms used by these studies presented lightly described contexts in which the shared goals of the conversation were not salient. In contexts like these, CWT predicts that literal contents will be interpreted as indeed yielding stronger commitments. Nevertheless, CWT also predicts that when contexts involve strong cooperation under specific goals non-literal contents can indeed generate testimonial

commitments. So far, however, no experimental paradigm has included a cooperation variable. For now, suffice it to say that the differential in accountability judgments is not necessarily counterevidence against CWT.

6. Conceptual Adequacy

I argued in §2 that the apparent dynamics of testimonial commitments present an extensional problem for conventionalist views of the communicative dimension of testimony. The epistemic commitments in **Charcoal** and **Charcoal*** are not captured by mere linguistic convention. Nevertheless, extensional adequacy is not all there is to theory building. Perhaps the cost of excluding pragmatically rich communication from testimony is worthwhile if there is an overwhelming conceptual link between convention and epistemic commitment.

In this section, I consider what this link could be. Ultimately, it is not clear what about convention gives it the exclusive privilege to constrain testimonial communication. In contrast, CWT offers a principled explanation of why testimony entails epistemic commitment.

In defending literalism about testimony, Fricker and Goldberg suggest that one of the main reasons we communicate beyond literal contents is to preserve deniability and remain able to deflect the kind of accountability that comes with testimony. To be sure, this is the case with our superintendent. His insinuation allows him to deflect epistemic deferrals and claims of accountability. So, one may suggest that conventions are especially suited to constrain testimony because they block this kind of deniability.

Nevertheless, the possibility of the kind of communicative exploitation just described extends to semantic conventions.²⁹ As the argument above suggests, even if the semantic components of an uttered sentence are clearly laid out in speech, speakers may concoct scenarios in which they gain plausible deniability by exploiting the context-sensitivity of literal contents. Speakers may, for instance, exploit ambiguous expressions (“oh sorry by ‘Geneva’ I meant Geneva, New York, not Geneva, Illinois”), pronouns (“oh, by ‘he’ I meant Jack, not Bert”), gradable adjectives (“oh, by ‘tall’ I meant tall for a Colombian”), and relational expressions (“oh, I meant I was ready to study, not ready for the exam”).

This problem extends to any wider communicative convention. Virtually all accounts of discourse conventions admit that a vast number of these conventions are ambiguous, and interlocutors must appeal to contextual features to disambiguate communicated contents.³⁰ Therefore, the use of discourse conventions—to whatever extent they are actually appealed in

²⁹ This point is also noted by Camp (2019), Hawthorne (2012), and Peet (2015).

³⁰ In general, Lepore and Stone (2015a) argue that what many theorists call pragmatic processes really amount to disambiguation among conventional alternatives. When it comes to the communicative dimension of testimony, this is precisely why convention alone in itself cannot explain testimonial commitments.

speech—is open to the kind of exploitation that convention was supposed to block. Hence, unless conventionalism is implausibly constrained to limit testimony to utterances that include explicit, in-sentence supplementation and disambiguation (“he is at the bank; and by ‘he’ I mean Jack Smith and by ‘bank’ I mean the financial institution in Linn Street”), defenders of the view must admit that the mere possibility of communicative exploitation is not enough reason to exclude non-literal contents from testimonial communication.

It might be tempting for conventionalists to argue that conventional communication may not always block the kind of exploitation enacted by our superintendent, but it is more epistemically stable than non-conventional communication. After all, by definition, conventions are procedures with a history of coordination success (see, for instance, Geurts 2018; Lepore & Stone 2015a; 2015b; Lewis 1975; 1980; Marmor 2009). Such history of stable behavioral patterns may give reason to audiences to think that when a convention is used, speakers mean the contents associated with it. Perhaps, then, conventions give audiences stronger inductive reasons to believe that speakers mean their associated contents.

At most, however, this is an argument for adopting conventionalism as a heuristic (and given my arguments above, not a great one.) It is not an explanation of why in principle the commitments that characterize testimony must be restricted to conventional contents. The conventionalist must explain not only why conventions can give us reasons to believe that speakers mean the associated contents but why the use of convention in particular is exclusively privileged in generating the epistemic commitments that characterize testimony.³¹

Inductive stability, however, does not fulfill this explanatory role. Jill may know that her neighbor Jack tends to turn the lights off when he leaves his house. But in no way can Jill hold Jack accountable for leading her to believe that he was home the day he forgot to turn off the lights. If all there is to say about the conceptual link between convention and testimonial commitment is that people tend to mean the contents associated with conventions, the case for a conventionalist account of the communicative dimension of testimony is not very strong.

In contrast, CWT offers a straightforward explanation of the link between cooperation and testimonial commitment. If CWT is correct, epistemically committing to a content follows as a means to achieve a shared goal. Therefore, epistemic failures in testimony are failures to comply with what speakers presented themselves as cooperatively committed. Given that communicating the truth of cooperatively warranted interpretations is required to fulfill shared goals, failing to have the appropriate epistemic warrant with respect to these

³¹ Could the conventionalist suggest that there is a further convention to hold speakers accountable for their testimony? Two considerations against this suggestion. First, this is a far cry from the plausibility and stability motivations that have motivated semantic conventionalism. Second, this would require the posit of a meta-grammar of accountability that, presumably, developed *after* linguistic conventions. CWT offers a more parsimonious explanation of why communication and accountability are deeply related.

contents entails failing to achieve said goals. In as much as failing to do what we commit to do with others entitles them to criticize and hold us accountable, failing to offer properly warranted testimony given a shared goal also entitles audiences to criticize and hold speakers accountable.³² Therefore, we can say that the specific contents derived from inferences whose crucial premise is the cooperative commitments of a speaker *inherit* the normative features of the commitment to the joint goal that warrants them.³³

Crucially, this strategy neatly explains why certain information, even if obtained from communication, does not carry the commitments that characterize testimony. As with our superintendent, sometimes, we can infer information given what others tell us through inferences that are not based on our shared goals. Given that their commitment to such shared goal is not what warrants the inferred content, speakers are not accountable for having communicated said content—certainly not in the way in which he would be had he testified to it. In our superintendent case, given his obvious goal of avoiding confession, he and his audience do not share the goal of exchanging information about bribing the commissioner. However his audience interprets the superintendent’s utterance, it will not be warranted by this goal and therefore will not inherit the epistemic commitments of testimony.

7. Conclusion

I have presented a distinctive account of the communicative dimension of testimony. CWT provides a principled criterion that individuates the contents for which testifying speakers are epistemically accountable. The practical normativity of cooperative action provides a context-sensitive account of the communicative mechanisms that serve a testimonial function in a given conversation. On the one hand, the cooperative criterion provided here is orthogonal and thus more comprehensive than accounts that appeal to conventions of speech. CWT in turn accommodates conventionalist intuitions by noting that the cooperative features of a conversation may warrant only minimal goal overlap and offers predictions of when testimony can ensue through non-conventional means. On the other hand, the cooperative criterion provides a principled explanation of why testimony as an epistemic source involves the kinds

³² This claim can have two readings. First, one may say that joint commitment *constitutively entails* entitlements of criticizing and holding accountable agents who fail to do what is required by the joint goal. Gilbert (2009) has an account of this sort. Second, one may say that joint commitments are merely concomitant with these entitlements. Bratman (2009; 2014) has a view of this sort.

³³ Testimony is often compared to promising (cf. Fricker 2006; Goldberg 2012; 2015; MacFarlane 2011), as it is seen as a speech act in which speakers willingly enact commitments for which they are responsible. Following that vein, CWT will be parallel to accounts of promising that explain the normativity of this speech act as generated by joint commitments in action. See de Kenessey (2020) for a promising account in this spirit and discussions on how it positively fares compared to conventional accounts.

of commitments that characterize it. These commitments are ultimately explained by the practical commitments to cooperate with an audience.

References

- Abusch, Dorit (2010). Presupposition Triggering from Alternatives. *Journal of Semantics* 27(1), 37–80.
- Adler, Jonathan E. (2002). *Beliefs Own Ethics*. MIT Press.
- Asher, Nicholas and Alex Lascarides (2005). *Logics of Conversation*. Cambridge University Press.
- Asher, Nicholas and Alex Lascarides (2013). Strategic Conversation. *Semantics and Pragmatics*, 6(2), 1–62.
- Bach, Kent (1997). The Semantics-Pragmatics Distinction: What It Is and Why It Matters. In Eckard Rolf (Ed.), *Pragmatik: Implikaturen Und Sprechakte, Linguistische Berichte* (33–50). VS Verlag für Sozialwissenschaften.
- Bach, Kent (2005). Context Ex Machina. In Zoltan Gendler Szabo (Ed.), *Semantics Versus Pragmatics*, (15–44). Oxford University Press.
- Bach, Kent and Robert M. Harnish (1979). *Linguistic Communication and Speech Acts*. MIT Press.
- Beaver, David and Jason Stanley (2018). Toward a Non-Ideal Philosophy of Language. *Graduate Faculty Philosophy Journal*, 39(2), 503–47.
- Borg, Emma (2010). Minimalism and the Content of the Lexicon. In Erich Rast and Luiz Carlos Baptista (Eds.), *Meaning and Context* (51–77). Peter Lang.
- Borg, Emma (2019). Explanatory Roles for Minimal Content. *Noûs* 53(3): 513–39.
- Bratman, Michael E. (1992). Shared Cooperative Activity. *The Philosophical Review*, 101(2), 327–41.
- Bratman, Michael E. (1993). Shared Intention. *Ethics*, 104(1), 97–113.
- Bratman, Michael E (2009). Modest Sociality and the Distinctiveness of Intention. *Philosophical Studies* 144(1), 149–65.
- Bratman, Michael E. (2014). *Shared Agency Shared Agency: A Planning Theory of Acting Together*. Oxford University Press.
- Burge, Tyler (2013). *Cognition Through Understanding: Self-Knowledge, Interlocution, Reasoning, Reflection: Philosophical Essays, Volume 3*. Oxford University Press.
- Camp, Elisabeth (2019). Insinuation, Common Ground, and the Conversational Record. In Daniel Fogal and Daniel W. Harris. *New Work On Speech Acts* (40-66). Oxford University Press.
- Cappelen, Herman, and Ernest Lepore (2005). *Insensitive Semantics: A Defense of Semantic Minimalism and Speech Act Pluralism*. Wiley-Blackwell.

- Carston, Robyn (1999). The Semantics/Pragmatics Distinction: A View From Relevance Theory. In Ken Turner (Ed.) *The Semantics/Pragmatics Interface From Different Points of View* (85-125). Elsevier.
- Coady, C. A. J. (2000). *Testimony: A Philosophical Study*. Oxford University Press.
- Cohen, L. Jonathan (1982). What Is Necessary for Testimonial Corroboration? *The British Journal for the Philosophy of Science*, 33(2), 161–64.
- Cohen, Philip and Hector Levesque (1988). Rational Interaction as the Basis for Communication (Technical Note 433) SRI International. Menlo Park, California.
- Cohen, Philip R. and Hector J. Levesque (1991). Teamwork. *Noûs*, 25(4), 487–512.
- Cohen, Philip R. and C. Raymond Perrault (1979). Elements of a Plan-Based Theory of Speech Acts. *Cognitive Science*, 3(3), 177–212.
- Douven, Igor (2006). Assertion, Knowledge, and Rational Credibility. *The Philosophical Review*, 115(4), 449–85.
- Donnellan, Keith S. (1968). Putting Humpty Dumpty Together Again. *The Philosophical Review*, 77(2), 203–15.
- Fricker, Elizabeth (1994). Against Gullibility. In Bimal Krishna Matilal and Arindam Chakrabarti (Eds.), *Knowing from Words: Western and Indian Philosophical Analysis of Understanding and Testimony* (125–61). Synthese Library. Springer Netherlands.
- Fricker, Elizabeth (2006). Second-Hand Knowledge. *Philosophy and Phenomenological Research*, 73(3), 592–618.
- Fricker, Elizabeth (2012). Stating and Insinuating. *Aristotelian Society Supplementary Volume*, 86(1), 61–94.
- Geurts, Bart (2018). Convention and Common Ground. *Mind & Language*, 33(2), 115–29.
- Gilbert, Margaret (2009). Shared Intention and Personal Intentions. *Philosophical Studies*, 144(1), 167–87.
- Gilbert, Margaret, and Maura Priest (2013). Conversation and Collective Belief. In Alessandro Capone, Franco Lo Piparo, and Marco Carapezza (Eds.) *Perspectives on Pragmatics and Philosophy* (1–34). Cham: Springer International Publishing.
- Goldberg, Sanford C. (2007). *Anti-Individualism: Mind and Language, Knowledge and Justification*. Cambridge University Press.
- Goldberg, Sanford C. (2008). The Epistemic Utility of What Is Said. In Robert Stainton and Christopher Viger (Eds.), *Compositionality, Context, and Semantic Values: Essays in Honor of Ernie Lepore* (61–78). Springer.
- Goldberg, Sanford C. (2012). *Relying on Others: An Essay in Epistemology* (reprint ed.). Oxford University Press.
- Goldberg, Sanford C. (2015). *Assertion: On the Philosophical Significance of Assertoric Speech*. Oxford University Press.
- Greco, John (2020). *The Transmission of Knowledge*. Cambridge University Press.
- Grice, Herbert Paul (1967). Logic and Conversation. In *Studies in the Way of Words* (41–58). Harvard University Press.

- Harris, Daniel W. (2019). We Talk to People, Not Contexts. *Philosophical Studies*. 177 (1), 2713–2733
- Harman, Gilbert (1990). *Reasoning, Meaning, and Mind*. Oxford University Press.
- Hawthorne, John (2012). Some Comments on Fricker’s ‘Stating and Insinuating.’ *Aristotelian Society Supplementary Volume*, 86(1), 95–108.
- Hanna, Joy, and Michael Tanenhaus. (2005). “The Use of Perspective during Referential Interpretation.” In John C. Trueswell and Michael K. Tanenhaus (Eds.) *Approaches to world-situated language use: Bridging the language-as-product and language-as-action tradition* (133-152). MIT Press.
- Hobbs, Jerry R. (1985). *On the Coherence and Structure of Discourse*. CSLI Publications.
- Hobbs, Jerry R., Mark E. Stickel, Douglas E. Appelt, and Paul Martin (1993). Interpretation as Abduction. *Artificial Intelligence* 63(1): 69–142.
- King, Jeffrey C. (2014). The Metasemantics of Contextual Sensitivity. In Brett Sherman and Alexis Burgess (Eds) *Metasemantics: New Essays on the Foundations of Meaning*. Oxford University Press, 97–118.
- Kenessey, Brendan de (2020). Promises as Proposals in Joint Practical Deliberation. *Noûs*, 54(1), 204–32.
- Kölbel, Max. 2011. Conversational Score, Assertion and Testimony. In Jessica Brown and Herman Cappelen (Eds.) *Assertion: New Philosophical Essays* (49-77). Oxford University Press.
- Korta, Kepa, and John Perry (2006). Varieties of Minimalist Semantics. *Philosophy and Phenomenological Research* 73(2): 451–59.
- Kvanvig, Jonathan L. (2009). Assertion, Knowledge, and Lotteries. In Duncan Pritchard and Patrick Greenough (Ed.), *Williamson on Knowledge* (140–60). Oxford University Press.
- Lackey, Jennifer (2008). *Learning From Words: Testimony as a Source of Knowledge*. Oxford University Press.
- Lepore, Ernest and Matthew Stone (2015a). David Lewis on Convention. In Barry Loewer and Jonathan Schaffer (Eds.) *A Companion to David Lewis*, (313–27). John Wiley & Sons.
- Lepore, Ernest and Matthew Stone (2015b). *Imagination and Convention: Distinguishing Grammar and Inference in Language*. Oxford University Press.
- Lewis, David (1975). Languages and Language. In Keith Gunderson (Ed.), *Minnesota Studies in the Philosophy of Science* (3–35). University of Minnesota Press.
- Lewis, David (1980). *Convention*. Blackwell.
- MacFarlane, John (2011). What Is Assertion. In Jessica Brown and Herman Cappelen (Eds.), *Assertion* (79–96). Oxford University Press.
- Marmor, Andrei (2009). *Social Conventions: From Language to Law*. Princeton University Press.

- Mazzarella, Diana, Robert Reinecke, Ira Noveck, and Hugo Mercier (2018). Saying, Presupposing and Implicating: How Pragmatics Modulates Commitment. *Journal of Pragmatics* 133: 15–27.
- MacKay, Alfred F (1968). Mr. Donnellan and Humpty Dumpty on Referring. *Philosophical Review* 77(2): 197–202.
- McKinnon, R. (2015). *The Norms of Assertion: Truth, Lies, and Warrant*. Palgrave Macmillan.
- McMyler, Benjamin (2011). *Testimony, Trust, and Authority*. Oxford University Press.
- Moran, Richard (2018). *The Exchange of Words: Speech, Testimony, and Intersubjectivity*. Oxford University Press.
- Moran, Richard A. (2013). I: Testimony, Illocution and the Second Person. *Aristotelian Society Supplementary Volume*. 87(1), 115-135.
- Neale, Stephen (2004). This, That, and the Other. In Anne Bezuidenhout and Marga Reimer (Eds.), *Descriptions and Beyond* (68–182). Oxford University Press.
- Owens, David. 2006. Testimony and Assertion. *Philosophical Studies* 130(1): 105–29.
- Owens, David. 2017. Human Testimony. In *Normativity and Control*, Oxford: Oxford University Press.
- Peet, Andrew (2015). Testimony, Pragmatics, and Plausible Deniability. *Episteme*, 12(1), 29–51.
- Peet, Andrew (2019). Knowledge-Yielding Communication. *Philosophical Studies*, 176(12), 3303–27.
- Pinker, Steven, Martin A. Nowak, and James J. Lee (2008). The Logic of Indirect Speech. *Proceedings of the National Academy of Sciences* 105(3): 833–38.
- Schiffer, Stephen (2003). *The Things We Mean*. Oxford University Press.
- Roberts, Craige (2012). Information Structure in Discourse: Towards an Integrated Formal Theory of Pragmatics. *Semantics and Pragmatics*. 5(6), 1-69.
- Roberts, Craige (2018). *New Work on Speech Acts Speech Acts in Discourse Context*. Oxford University Press.
- Saul, Jennifer Mather (2012). *Lying, Misleading, and What is Said Lying, Misleading, and What Is Said: An Exploration in Philosophy of Language and in Ethics*. Oxford University Press.
- Saul, Jennifer (2018). *New Work on Speech Acts Dogwhistles, Political Manipulation, and Philosophy of Language*. Oxford University Press.
- Schiffer, Stephen R. (1972). *Meaning*. Clarendon Press.
- Schiffer, Stephen R. (2003). *The Things We Mean*. Oxford University Press.
- Searle, John R. (1965). What Is a Speech Act? In Max Black (Ed.), *Philosophy in America* (221–39). Cornell University Press.

- Simons, Mandy. 2010. A Gricean View on Intrusive Implicatures. In Klaus Petrus (Ed.) *Meaning and Analysis: New Essays on Grice* (138–69). Palgrave Macmillan
- Simons, Mandy. 2018. Convention, Intention, and the Conversational Record. In Gherard Preyer (Ed.) *Beyond Semantics and Pragmatics* (284-302). Oxford University Press.
- Stainton, Robert J. 2016. Full-On Stating. *Mind & Language* 31(4): 395–413.
- Stalnaker, Robert (2014). *Context*. Oxford University Press.
- Stanley, Jason (2008). Knowledge and Certainty. *Philosophical Issues*, 18(1), 35–57.
- Sternau, Marit, Mira Ariel, Rachel Giora, and Ofer Fein (2015). Levels of Interpretation: New Tools for Characterizing Intended Meanings. *Journal of Pragmatics*, 84(1), 86-101
- Stokke, Andreas and Anders J. Schoubye (2016). What Is Said? *Noûs*, 50(4), 759–93.
- Thomason, Richmond (2003). Accommodation, Meaning, and Implicature: Interdisciplinary Foundations for Pragmatics. In Philip R. Cohen, Jerry Morgan, and Martha E. Pollack (Eds.) *Intentions in Communication*. (235-363). MIT Press.
- Tuomela, Raimo (2006). Joint Intention, We-Mode and I-Mode. *Midwest Studies In Philosophy*, 30(1), 35–58.
- Welker, Katherine. 1994. Plans in the common ground: toward a generative account of conversational implicature. Columbus. (PhD thesis.) Ohio State University
- Williamson, Timothy (1996). Knowing and Asserting. *The Philosophical Review*, 105(4), 489–523.

CHAPTER 2

Committing to a Fault: The Ethics of Communicative Commitment

If I were to deceive you into falsely believing that Ana bribed the police, saying ‘she bribed them’ seems worse than merely saying ‘she solved it like a businesswoman.’ This suggests that there is a normative difference between declarative speech acts that commit speakers to a content and declarative speech acts that do not. I argue that a vastly popular rendering of this normative differential is misguided. Call it Literalism: the idea that the normative distinction between committal and noncommittal speech traces the moral distinction between literal and nonliteral speech. Against Literalism, I suggest that there is a pervasive kind of nonliteral speech that systematically carries the same commitments as literal speech: Nonliteral cooperative signals. Hence, I argue, we should endorse Cooperativism, which holds that the normative difference between committal and noncommittal speech is best explained by appealing to the distinction between contents that are required and contents that are merely allowed as cooperative moves in a given conversation. Cooperativism, in turn, i) accommodates context-dependent cases of literal and nonliteral commitment, ii) better captures the moral principles that seemed to lend support to Literalism, and iii) offers a principled pathway to theorize about the role of social phenomena (like power, institutions, and social norms) in the ethics of speech.

1. Overview

Some declarative speech acts commit the speaker to a content. Others not so much. Say that Miguel is having a casual conversation with his office mates about that time Ana was detained by the police and suddenly let go. He wants to deceive his audience into believing that she bribed the police. Consider two possible utterances:

- 1) She bribed the police.
- 2) She solved it like a businesswoman.

Of course, since Miguel is deceiving his audience, both utterances are at least *prima facie* wrong. However, at least intuitively, Miguel’s deception in (1) seems *worse* than his deception in (2). Having uttered (2), he is able to validly deflect at least some criticisms that seem entirely appropriate had he uttered (1)—“I never told you that she bribed the police!”. Therefore, there seems to be a normative distinction between deception through declarative speech acts that commit the speaker to a content and deception through declarative speech acts that do not.

Note also that this intuition can be generalized. Say that Miguel does not have the intention to deceive, but is just irresponsibly guessing that Ana bribed the police. It still seems that his uttering (1) is worse than his uttering (2). Again, in uttering (2), he seems to be able to deflect criticisms that are entirely appropriate had he uttered (1). Therefore, we can posit an at least plausible claim:

Normative Differential: There is a normatively significant difference between committal and noncommittal communication.

That said, my argument here will be guided by two questions: A) How should we characterize the distinction between committal and noncommittal communication (if there is one at all)? And B) Why or in virtue of what is this distinction normatively significant (if it is at all)?

As a foil to the view I develop here, I focus on a vastly popular way of answering these questions.

Literalism: The distinction between committal and noncommittal communication traces the moral distinction between communicating literal contents and communicating nonliteral contents.

The Literalist holds that the distinction between committal and noncommittal communication is explained by appeal to the literal use of semantic conventions in conversation. The relevant difference between (1) and (2) is that the content [Ana bribed the police] is encoded in the literal content of (1), while it is not encoded in the literal content of (2). Further, for the Literalist, the normativity at stake is moral. There is something about the use of semantic conventions in conversation such that intentionally defective literal communication breaches moral norms not breached by intentionally defective nonliteral communication.

I argue against two strains of Literalism: views that appeal to moral norms that govern the relationship between a speaker and her audience¹ and views that appeal to moral norms that govern the relationship between a speaker and her linguistic community.² The literal use of semantic conventions in communication, I argue, is a special case of a general speech act-kind: Cooperative signaling. However, cooperative signals can be performed nonliterally. And, I argue, if the moral norms to which Speaker-Audience and Speaker-Community strains of literalism govern literal cooperative signals, these norms govern nonliteral cooperative signals as well. Hence, literalist arguments fail to morally adjudicate between literal and nonliteral communication.

Cooperative signaling, in turn, serves as a springboard for my positive view. Contrary to other opponents of Literalism, I argue that there is still a normatively significant difference between committal and noncommittal communication.³ However, this distinction should be

¹ Key examples are Adler (1997; 2018), Chisholm & Feehan, (1977) Gormally, (1991), MacIntyre (2006), Mahon (2003), Strudler (2010), Timmermann & Viebahn (2021). Classic versions of it can be found in Augustine (2002[395]a, 2002[395]b).

² This strain of Literalism has been recently defended by Berstler (2019). But see also Owens (2022) for a similar argument.

³ Compare to Saul (2012) and Williams (2004).

explained in terms of the rationality of non-language-specific cooperative action. More specifically, I endorse the following claim:

Cooperativism: The distinction between committal and noncommittal speech corresponds to the distinction between contents that are *required* and contents that are *merely allowed* as cooperative interpretations of an utterance in a given conversation.

The difference between (1) and (2) is that interpreting the content [Ana bribed the police] is required to make sense of (1) as a cooperative move in the conversation. In contrast, although the same content is allowed as a cooperative interpretation of (2), there are other alternative contents that stand on equal interpretive grounds as cooperative moves in the conversation. Crucially, this distinction is orthogonal to whether cooperatively required contents are communicated literally or nonliterally.

Under cooperativism, the normative distinction between committal and noncommittal communication is grounded on the instrumental normativity of cooperative or joint action. Intentionally communicating a defective cooperatively required content, we will see, is a constitutive betrayal of the cooperative commitments that enabled it in the first place. This kind of failure does not obtain with deceptive merely allowed contents. However, whether this normative distinction amounts to a moral distinction will depend on more general commitments about the ethics of joint action. Although these concerns are beyond the scope of this paper, I conclude by arguing that even if this normative distinction is not moral itself, Cooperativism yields novel insights for understanding the ethics of communication.

In the next section I say more about Literalism and its Speaker-Audience (§2.1) and Speaker-Community (§2.2) strains. In (§3), I offer an argument for individuating the general speech-act kind of cooperative signaling and how it is plausibly governed by the moral norms appealed to by the two strains of literalism, no matter whether this act is instanced literally or nonliterally. In (§4) I present cooperativism. I introduce the normatively significant distinction between contents that are required (§4.1) and contents that are merely allowed (§4.2). In (§5) I conclude by explaining how cooperativism illuminates the ethics of speech by offering a framework to understand the role of higher-order social phenomena in the ethics of communication.

2. Literalisms

Recall our initial questions:

- A) How should we characterize the distinction between committal and noncommittal speech (if there is one at all)?
- B) Why, or in virtue of what, is this distinction normatively significant (if it is at all)?

In this section I present a widely popular answer to these questions: Literalism. For starters, I say something about the Literalist answer to (A); then, I turn to two different answers to (B).

For our purposes, literal contents can be characterized as the contents determined by the semantic components of an uttered sentence, given their syntactic structure, and after the contextual supplementation required to express a truth-evaluable proposition. They correspond (or are close to), for instance, what has been called semantic contents⁴, minimal contents⁵, and what is strictly said.⁶ Importantly, literal contents are context dependent to a significant degree and allow for the supplementation of some expressions such as ambiguous terms ('There are great restaurants in Geneva [New York, not Illinois]'), pronouns ('he [John Smith] is an engineer'), relative adjectives ('Andrés is tall [for a Colombian]') and relational expressions ('Pam is ready [for the exam]'). However, they exclude any additional logical, presuppositional, and conversational implications that go beyond this kind of supplementation.

To be sure, what exactly is the semantically determined literal content of an uttered sentence—and whether this is even a plausible construct—is a contentious issue.⁷ Moreover, intuitions about what counts as the literal content of a number of expressions are notably muddy.⁸ Nevertheless, it seems intuitive that there is something like the literal content of an uttered sentence. So, for our purposes, let us assume that most (if not all) uttered sentences have literal contents.

I focus on principled reasons in favor of literalism. For this, it is important to note that literal contents are conventionally determined. They are determined and constrained by the linguistic/semantic conventions accepted by a given community.⁹ The fact that speaker and audience know and accept the same linguistic conventions (of Spanish, English, or even Klingon) allows them to issue and interpret the literal contents of their uttered sentences. Nonliteral contents, on the other hand, require nonconventional determinants to be issued and interpreted, such as contextual clues concerning speaker intentions. In turn, the Literalist holds, there is something about semantic convention that renders literal communication

⁴ Cf, Saul (2012) Bach (2006)

⁵ Cf., Borg, 2010, 2019, Cappelen & Lepore, 2005

⁶ Cf, Goldberg (2008)

⁷ See, for instance, Bach (1997; 2005), Borg (2007), Carston (1999), Grice (1968), Korta & Perry (2022), Perry (2007), Recanati (2003; 2003), Stanley & Szabó (2000), Stokke & Schoubye, (2016)

⁸ Some empirical work has been done concerning this threshold, see Mazzarella, et al. (2018), Reins & Wiegmann (2021). Results do not seem to support Literalism as presented here. However, see Michaelson (2013) for an optimistic view on the role of these intuitions.

⁹ Following the literalist literature, I focus on semantic conventions. Discourse conventions are arguably wider than that (e.g., Lepore & Stone 2015b). One may argue for a wider version of literalism on these grounds. Since this moral conventionalist view has been defended in the literature and my argument below arguably extends it, I bracket this issue.

subject to distinctive moral norms that do not govern nonliteral communication. In the next two subsections, I discuss two kinds of argument for this claim.

Before, however, two clarificatory notes. First, I am concerned here with Literalism as a moral thesis, i.e., as a view on the moral differential between literal and nonliteral communication. There are versions of Literalism that hold that there is an *epistemic* difference between literal and nonliteral communication.¹⁰ Although I believe epistemic Literalism is prey to similar challenges to the ones I present here, showing why falls outside the scope of this paper.¹¹

Second, arguments in favor of moral Literalism (henceforth just Literalism) are mostly offered to characterize the moral distinction between lying and misleading. For our purposes and in a broad stroke, lying amounts to deceiving through literal means and misleading amounts to deceiving through nonliteral means.¹² My aim here is to argue that the distinction between literal and nonliteral communication, and thus the distinction between lying and misleading, is not morally significant. However, I hold that there still is a normatively significant distinction between committal and noncommittal communication—it is just not grounded in semantic conventions. This may motivate shifting the distinction between lying and misleading, so that lying can be done nonliterally. Alternatively, it may be better to preserve the lying/misleading distinction along the literal/nonliteral distinction and abandon the thesis that it is morally relevant. Since I am not prepared to die on either of these taxonomical hills, I bracket the use of the terms ‘lying’ and ‘misleading’ and talk about literal and nonliteral deception.

2.1 *Speaker-Audience Literalism*

Classic arguments in favor of the moral difference between deceiving literally and deceiving nonliterally appeal to considerations of trust and autonomy concerning the relationship between a speaker and her audience. Something about deception through literal means betrays the audience in a way that deception through nonliteral means does not.

A crucial claim in these arguments is that literal deception *directly* betrays the faith of an audience, whereas nonliteral deception does so *indirectly*. As the argument goes, all that

¹⁰ A number of philosophers have argued or suggested that only literal communication can ground the commitments necessary for testimonial communication, e.g., Burge (1997), Fricker (2012), Goldberg (2007, 2008), MacFarlane (2011), Owens (2017).

¹¹ See Hawthorne (2012), Peet (2015, 2016), and Vesga (forthcoming) for considerations against epistemic Literalism.

¹² Stokke (2018) offers a more subtle and altogether plausible distinction between lying and misleading. Lying is a specific kind of deception that requires saying something false, where saying is constrained by semantically determined minimal contents close to what I have called the literal contents of our words. In turn, the contextual supplementation and determination of these contents is guided by Questions Under Discussion, which allows for a wider conception of lying that it is usually allowed in the literature. Notably, Stokke does not argue for the normative distinction between lying and misleading. So, for all we have both said, our arguments are compatible.

needs to happen for Miguel's deception in (1) to be successful is that he utters his words and his audience interprets them. In contrast, if (2) is to be successfully deceptive, Miguel's audience needs not only to decode the literal meaning of his words; they need to engage in an extra inferential, non-conventional process that leads them to believe that Ana bribed the police. This extra inferential step is supposed to yield a moral distinction between literal and nonliteral communication.

To illuminate the spirit of the argument, consider a Kantian scenario that is often appealed to in the literature.¹³ Lucía wants Rajiv to falsely believe she is leaving on vacation. So she packs her bags knowing that Rajiv will see them later and infer that she is planning to take a trip. Say that Rajiv actually sees the bags and indeed forms the false belief. It seems intuitive that this kind of deception is better than Lucía literally saying to Rajiv, 'I'm going on vacation.' This is because in the packing case, Rajiv's inference depends on his own epistemic capacities, of which he has full control. In contrast, in the literal deception case, Rajiv's given up his full epistemic control and formed his belief because Lucía assured him that he can put his trust in her. In literally deceiving and not in packing, Lucía has taken responsibility for Rajiv's belief and thus directly betrayed his trust.

Nonliteral communication, the argument goes, is parallel to the packing case. Miguel's audience must engage in an autonomous inferential process to arrive at the conclusion that Ana bribed the police. As Adler puts it, "[t]he underlying idea is, presumably, that each individual is a rational, autonomous being and so fully responsible for the inference he draws, just as he is for his acts." (Adler, 1997). Given the differences in the interpretive processes of literal and nonliteral communication, literal deception amounts to a direct betrayal of the audience's trust, while this is not the case with nonliteral deception.¹⁴

Why exactly this inferential difference amounts to a moral difference varies in the literature. For our purposes, consider three candidate moral norms that could render this alleged extra inferential step morally significant.

(A): Doing a wrong is worse than allowing a wrong.

(B): Blame is distributed among those who are responsible for a wrong.

(C): Actions sufficient for a wrong are worse than actions insufficient for a wrong.

¹³ See (Kant, 1996a, 1996b). Variants of this scenario are used, for instance, in Shiffrin (2016), Chisolm & Feehan (1997), Adler (2013) to defend the moral difference between literal and nonliteral deception.

¹⁴ Timmermann & Viebahn (2021) offer a different kind of argument. They appeal to Brandomian accounts of assertion, in which to assert is to enact social commitments. So, given that assertions are literally constrained, *ipso facto* they exclusively carry said commitments—as opposed to implicatures, for instance. While I am sympathetic to Brandomian accounts, they often lack an explanation of why are assertions literally constrained and why do they carry their distinctive commitments. If my argument below is right, the reasons why assertions carry commitments are the same reasons why a distinctive kind of non-literal speech carries these same commitments.

Under (A), the moral difference stems from the notion that speakers directly inflict, so to speak, a false belief on their audience. In contrast, misleading speakers merely allow their audience to arrive at the false belief on their own (Chisholm & Feehan, 1977; Adler, 2013). Under (B), one may admit that misleading speakers are responsible for deceiving, but that audiences share some of that responsibility because they engaged in an autonomous inferential process over which they had control (McIntyre, 2006; Shiffrin, 2015; Green, 2001). Under (C), intentionally conveying a false literal content is sufficient for the lying speaker to wrong their audience. In contrast, the misleading speaker needs to convey a literal content *and* trust their audience to engage in an extra inferential process (Strudler, 2010).¹⁵

Different forms of (A)-(C) have been defended or presumed by ethicists beyond the lying-misleading debate. As such, they are controversial. My aim here is not to argue for or against them. Rather, I argue in §3 that, if they hold, these norms do not adjudicate between literal and nonliteral communication. That is, there is a kind of nonliteral communication that is systematically subject to the same kinds of considerations of trust and autonomy to which speaker-audience literalism appeals.

2.2 *Speaker-Community Literalism*

Dissatisfied with the arguments above, Sam Berstler (2019) argues that literal deception is morally worse than nonliteral deception given moral considerations about the speaker's relationship with her linguistic community.¹⁶ In Berstler's view, literal deception breaches a norm entailed by communal cooperative ventures:

Fair Play: When a number of people engage in a just, mutually advantageous cooperative venture, according to rules, and thus restrain their liberty in ways necessary to yield advantage for all, and when you have benefitted from these people's submission, you are obligated to conform to the rules (2019: 78).

When summoned in other domains, the normative grip of Fair Play generally stems from the facts that the relevant practice is maintained by the costly compliance of a community, and said compliance yields positive outcomes for the community (e.g., Arneson, 1982; Hart, 1955; Rawls, 1999; Simmons, 1979). These conditions generate moral reasons to comply with the rules of the practice even if no harm actually follows from token actions that breach them. Breaching these rules is unfair to other members of the community who bear the costs of maintaining the practice.¹⁷

¹⁵ See Saul (2012) and Berstler (2019) for a thorough case against these arguments.

¹⁶ See Owens (2022) for a similar argument, although not explicitly concerning the moral difference between literal and non-literal deception.

¹⁷ Some philosophers have posed strong challenges to Fair Play, e.g., Nozick (2013).

Berstler argues that Linguistic conventions are a perfect candidate to be the kind of rules subject to Fair Play. She appeals to what Lewis (1975, 1980) calls the conventional practice of *telling the truth in a language*. A language is a system of arbitrary conventional communicative procedures—using the word ‘chair’ to refer to a chair, using the word ‘I’ to refer to the speaker, etc. And, crucially, languages are effective because speakers comply with the further, general convention of appropriately using the procedures that comprise them—actually using the word ‘chair’ to refer to chairs, for instance. This general convention of using generally the language for the purposes for which it was developed is what Lewis calls the convention of telling the truth in a language.

With this in place, Berstler argues, the comprehensive and adequate use of a language is enabled and maintained because enough community members comply with its conventions and thus enact the sacrifices needed to continually use the same language (such as avoiding lies or idiosyncratic uses of conventions.) In turn, this means that recurrent violations will erode the effectiveness of this practice. If enough members of a community violated the practice of telling the truth in a language, said language would cease being a useful tool for communication.

Under the rather plausible assumption that the communal use of a language is an overall beneficial practice, the use of linguistic conventions in speech is subject to Fair Play. Speakers who violate the conventional practice of truthfulness in a given language exploit the fact that others are making the sacrifices necessary to maintain and comply with it. Literal deceivers like Miguel in (1) are freeriding the linguistic conventions for their own benefit. In contrast, pragmatically rich speech like the one exploited by Miguel in (2) has meaning in virtue of instance-specific, extra-conventional facts of the conversation. It does not depend on communal effort and is therefore not breaching Fair Play.

3. Cooperative Signaling

I begin elucidating my positive argument by suggesting how the Speaker-Audience and Speaker-Community strains of Literalism are misguided. In short, they arbitrarily focus on semantic convention, while missing a critical normative aspect of communication in general.

For starters, consider two variants of a familiar scenario. Felipe and Andrea are going to the movies in Felipe’s car. Felipe is from out of town, and it is clear that they need gas. Andrea utters:

- 3) You can get gas at the station down the road.
- 4) There’s a station down the road.

Andrea’s utterance (4) is a paradigmatic case of what I call here a cooperative implicature.¹⁸ It communicates that Felipe can get gas at the station down the road—even if said content is not encoded by the semantic features of the uttered sentence. Given their common knowledge that they are going to the movies and Andrea is motivated to achieve this goal, (4) communicates that Felipe can get gas at the station down the road. Otherwise, it would not be a relevant conversational move.

At least to my ear, Andrea is as committed to the content [Felipe can get gas at the station down the road] in (1) as she is in (2). To bolster this intuition, consider two tests. First, say that Felipe asks, “can I get gas at the station down the road?”. That question seems superfluous as a follow-up to both utterances. Given their context, Andrea would not have uttered (4) unless she wanted to communicate so. Second, say that they arrive at the gas station and there is no gas. It seems that Andrea’s appropriate excuse is epistemic (“I thought they had gas!”) and not linguistic (“I never told you they had gas!”). This strongly suggests that (4) has the function of adding the content [Felipe can get gas at the station down the road] to the common ground and thus commits Andrea to the cooperative implicature.¹⁹

These, however, are intuitions. Our focus here is on principled arguments for and, in this case, against the moral difference between (3) and (4). My argument begins by noting that there is a crucial commonality between literal communication as exemplified by (3) and cooperative implicatures as exemplified by (2).

To see this commonality, let us appeal to a classic and widely accepted account of linguistic convention developed by Lewis.²⁰ In social life, agents often face coordination problems. They mutually acknowledge that their motivations are aligned to achieve the same world-state and to act mutually and responsively to do so. When these problems occur recurrently and on a communal scale, agents develop procedures to solve them effectively and efficiently. Drivers, for instance, often face the coordination problem of not crashing. Hence, communities develop arbitrary procedures to solve this recurrent problem, such as driving on one side of the road, for instance. These arbitrary procedures are conventions: general, mutually accepted action scripts meant for solving recurrent coordination problems.

Languages are conventions in this sense. Agents often encounter the coordination problem of communicating about the world—of providing each other with information about how the world is. Given this recurrent problem, communities develop arbitrary

¹⁸ ‘Implicature’ has come to have different meanings in the literature. By using ‘cooperative implicature’ I mean to mark the nonliteral contents that are communicated in virtue of their being cooperatively required as conversational moves; see Grice (1968). As we will see, there are other means of nonliteral communication whose effectiveness does not hinge on cooperative interpretations.

¹⁹ This considerations run against the considerations offered by Timmermann & Viebahn (2021).

²⁰ See too Geurts, 2018, Lepore & Stone, 2015a, 2015b, and Marmor, 2009. See Millikan (1987) for an alternative rendering of convention. *Mutatis mutandis*, my arguments should hold under her construal as well.

procedures to solve it. English speakers have the convention of using the word ‘chair’ to refer to chairs and apostrophes to denote possession. Spanish speakers, respectively, use the words ‘silla’ and ‘de’ for these same purposes. Languages (their declarative dimensions, at least) are conventional systems of communicative procedures developed to solve the coordination problem of informing and being informed about the world.

Importantly, *at a given instance* a convention is effective only if it is contextualized as aimed at solving the coordination problem it was developed to solve. In a demolition derby or a NASCAR race, the convention of driving on one side of the road is not effective because agents are not engaged in a coordination problem that contextualizes said action as a useful practice. Likewise, at a given instance, linguistic conventions fulfill their communicative purpose when they are contextualized as aimed at sharing information as an at least partial solution to a coordination problem.²¹

This means that the Speaker-Audience Literalist is wrong in thinking that all that is involved in literal communication is the issuing of a semantically involved signal and the interpretation of that signal’s conventional meaning. For an utterance to actually communicate its literal content, it must be embedded in a context that warrants said interpretation as a communicative move in conversation. This is what distinguishes literal communication from irony, sarcasm, pretense, and—perhaps more dramatically—sleep-talking. Say that Andrea utters (3) in her sleep. Felipe may well understand its meaning; he can even infer that he can get gas at the station down the road. Yet by no means Andrea has committed to said content. More schematically, at a given instance, an utterance communicates its literal content whenever said content is contextualized as a solution to a coordination problem—i.e., as a declarative conversational move aimed at advancing the interlocutors’ shared interests in solving the coordination problem of informing and being informed about the world.

Crucially, sometimes, the contextual features of a communicative engagement call for different, nonconventional interpretations to make sense of an utterance’s informative role in solving a coordination problem. The fact that Andrea has presented herself as interested in advancing the mutual interest of going to the movies warrants the interpretation of her utterance as a cooperative move towards this goal, which contextualizes said utterance as communicating that Felipe can get gas at the station down the road. Alternative interpretations—including the literal interpretation of (4)—would have rendered her

²¹ In terms of Thomas Schelling’s (1960) defining work, linguistic conventions are effective at a given instance because the fact that a sound has been conventionally associated with a content is recognized as the *focal point* that allows speaker and audience to coordinate in communicating about the world.

utterance irrelevant as a move directed towards solving their coordination problem.²² Note the similarity with literal communication. At a given instance, an utterance communicates a cooperative implicature whenever said implicature is contextualized as a solution to a coordination problem—i.e., as a declarative conversational move aimed at advancing the interlocutors in solving the coordination problem of informing and being informed about the world.

This gives us reason to think literal communication and cooperative implicatures are two special cases of a more general kind of action: issuing a signal that is interpreted as solving a coordination problem given the interlocutors' mutually acknowledged commitment to advance their shared interests. The fact that Felipe and Andrea have mutually acknowledged their shared interest in getting to the movies contextualizes both (3) and (4) as conversational moves aimed at communicating that Felipe can get gas at the station down the road. The literal interpretation of (3) and the nonliteral interpretation of (4) are warranted in virtue of the fact that they are solutions (or partial solutions) to coordination problems, given the mutually acknowledged interest in communicating about the world. I offer a more technical characterization of this act in §4. But for now, call this act cooperative signaling.

Importantly, this is one way to understand Lewis's distinction between telling the truth *in a language* and telling the truth *simpliciter*. The practice of telling the truth in a language amounts to using a given set of conventions (a language) to solve the coordination problem of informing and being informed about the world. Telling the truth *simpliciter* is the general practice of cooperatively issuing a signal that solves the coordination problem of informing and being informed about the world, no matter the means by which it is done. Hence, telling the truth in a language is a way of telling the truth *simpliciter*.

And, critically, telling the truth *simpliciter* is not a convention.²³ Rather, it is the name of the game; it is the aim that motivates communicative coordination between interlocutors. In our parallel example, the act of not crashing is not a convention in itself but the aim that motivates coordination. Of course, driving on one side of the road is an effective way to achieve this aim; but there are other ways to avoid crashing. Likewise, one may tell the truth *simpliciter* by telling the truth in a language or by issuing nonliteral cooperative implicatures. In both cases, speakers issue signals that, given the mutually acknowledged aims of the conversation and salient features of the context, are to be interpreted as communicating about the world if they are to be a rational move within the coordination engagement in which they are embedded.

²² Again, in Thomas Schelling's terms, speakers can communicate cooperative implicatures because the context, together with the acknowledged aims of the conversation, shift the *focal point* of the conversation from a literal interpretation to a nonliteral one. A literal interpretation is no longer the salient solution to the coordination problem.

²³ Cf. Lewis 1975, p. 184.

We are now in a position to see where the speaker-audience and speaker-community versions of Literalism go astray. In short, there is a systematic range of nonliteral communication that is arguably governed by the same moral norms that govern literal communication: nonliteral cooperative signals, mainly represented by what I have called here cooperative implicatures. Again, my aim here is not to argue for or against these norms; rather, I suggest that, even if these norms hold, they do not adjudicate between literal and nonliteral communication.

3.1 Against Speaker-Audience Literalism

Say that Andrea is a devious speaker. She knows that there is no gas at the station down the road but wants Felipe to form the false belief that there is. If so, in both (3) and (4) she has issued a deceptive cooperative signal. That is, she has presented herself in bad faith as motivated to solve a coordination problem—obtaining gas to go to the movies—and this contextualizes her utterances as communicating that Felipe can get gas at the station down the road. In both variants of the case, if her deception is to be effective, Andrea must present herself as cooperatively committed to said goal to even be able to issue the contents at stake.

Note that the parallel between Andrea’s nonliteral cooperative implicature in (4) and Lucía’s packing does not hold. Felipe’s interpretation of Andrea’s implicature is anchored in her public and active commitment to advance the shared goal that guides their conversation. Such a commitment is not a premise in Rajiv’s inference when he sees Lucía pack her bags. Lucía’s packing-based deception may very well be less wrong than overtly lying by saying “I’m leaving on vacation.” But by no means this entails that we should lump all nonliteral communication with Lucía’s packing-based deception.

More generally, cooperative signals (both literal and nonliteral) are enabled by a distinctive kind of non-language-specific shared agency that is only successful when each agent trusts that the other will do their part. To coordinate in bringing the couch downstairs, we have to acknowledge our mutual commitment to do so and act in a way that would be successful only if we both acted together. This means that, in coordinating with each other, agents surrender full control over the success of their actions and trust that the other will act accordingly to achieve a goal together. If only I do what is required to bring the couch together, we would fail to bring the couch downstairs *together*—even if I somehow manage to do it on my own. Therefore, intentionally failing to fulfill the cooperative commitment to coordinate amounts to intentionally breaching the trust that another agent has placed in us to do our part in acting together.

Presenting oneself as committed to this kind of joint agency—to doing something with another agent—enables us to contextualize our actions as cooperative signals (literal or not.) Intentionally deceiving with a cooperative signal thus amounts to intentionally failing to

fulfill the shared commitment that warrants the interpretation of one's utterances as cooperative signals. Andrea has actively presented herself (albeit in bad faith) as committed to a shared goal, which in turn warrants Felipe's interpretation in both (3) and (4) that he can get gas at the station down the road. If she is a devious speaker, she betrays the very same commitment that enabled her to issue her cooperative signal in both cases. There is no extra inference that serves as a moral difference-maker between Andrea's literal and nonliteral cooperative signal.

Therefore, if (A)-(C) above govern communication at all, they govern a systematic and pervasive kind of nonliteral communication, i.e., nonliteral cooperative signals. Given the context at stake, Andrea is not merely letting Felipe arrive at the conclusion that he can get gas at the station down the road. Rather, she is actively guiding him to that conclusion by way of presenting an utterance as advancing her cooperative commitment to get to the movies. Hence, there is no principled reason to think that responsibility distribution will systematically differentiate the use of literal speech and the use of pragmatically rich speech like cooperative implicature. Likewise, in both literal and pragmatically rich communication, the same act is sufficient for the wrong: presenting oneself as a cooperative interlocutor and intentionally yielding a false cooperative interpretation of one's words.

3.2 *Against Speaker-Community Literalism*

Recall our rendering of Lewis's distinction between telling the truth *simpliciter* and telling the truth in a language. Telling the truth *simpliciter* amounts to the act of sharing information to solve a coordination problem. In our terms, it amounts to issuing a truthful cooperative signal. Telling the truth in a language is one way of doing so; it amounts to the use of accepted linguistic conventions to tell the truth *simpliciter*. However, as we have seen, cooperative implicatures show that one can also tell the truth *simpliciter* with nonliteral conversational moves.

Hence, speakers who deceive through literal *and* nonliteral cooperative signals betray the aim of telling the truth *simpliciter*. They present themselves in bad faith as committed to advancing the shared goals that contextualize their utterances as aimed to tell the truth. The question, then, is whether the general practice of telling the truth *simpliciter* is subject to communal norms breached by both literal and nonliteral cooperative signals.

To begin, note that linguistic conventions need not be a communal matter. You and I may develop our own private language and use it to fulfill its intended purpose or to betray and deceive each other. In fact, new communicative conventions constantly emerge in local interactions between strangers, interlocutors who do not speak the same language, and children and their parents. Hence, if Berstler is right, literal communication is subject to communal norms to the extent to which linguistic conventions are maintained by and

through communal compliance (which is admittedly a rather significant extent.)²⁴ In parallel, communal norms arguably govern the general practice of telling the truth *simpliciter* to the extent to which it is maintained by and through communal compliance. And, crucially, this is very much the case.

Human cooperative communication is a communal pattern of behavior. When we communicate with others, we do so against a rich backdrop of previous cooperative interactions. This background of collective behavior enables communicative participation on the scale that distinguishes humans from other animals; it is why I can trust a stranger when they tell me that I can buy groceries in the downtown mall—whether they utter, “You can buy groceries in the downtown mall”, or “There’s a mall downtown”, or “The mall downtown.” This backdrop of sustained cooperation enables effective communication by reducing the costs of deliberating whether every person we talk to is properly motivated to coordinate with us about the world. Call this the *communal presumption of cooperation*. Under this presumption, agents can communicate widely and variably—literally and nonliterally—with agents with whom they have not had any previous communicative interactions.

This presumption, however, is defeasible by recurrent violations.²⁵ A community whose members constantly betray cooperative communicative commitments will progressively erode the possibility that cooperative communication can serve as a default. A sufficient number of betrayals will make it progressively harder to engage in effective communication supported by the presumption of cooperation. The communal scale that enables the widespread effectiveness of the presumption of cooperation is therefore maintained by the efforts of the members who comply with its standards of truthfulness, which will not be conventionalized rules, but constitutive requirements for large-scale

²⁴ From an evolutionary perspective, it is odd to think that languages are maintained by the communal efforts and sacrifices of a community. Languages are cultural products that coevolved with our social cognition (see Tomassello (2010); Henrich 2019; Christiansen & Chater 2022). They emerged because maintaining a communal repertoire of communicative procedures was advantageous to our ancestors, given our social and natural environments. This gives us reason to think that, as cultural products, languages are selected and maintained by evolutionary mechanisms, not by the efforts and sacrifices of a community. Freeriders, of which there have certainly been plenty in our species’ history, have done nothing to erode the presence and effectiveness of linguistic tools. In fact, there is reason to think that as long as language use is advantageous to our species as a whole, there will always be advantages for a minority of deceptive token assertions. However, the massive advantages of truth will ensure that this equilibrium is maintained. This is an argument against language being the kind of social practice subject to Fair Play. Since my aim is to show that, *if* Fair Play holds, it does not differentiate between literal and nonliteral communication I bracket this argument in what follows. Note, however that, my argument below is not subject to this kind of objection. Insofar as *fair* communal practices motivate communicative coordination, then Fair Play applies to cooperative signals (be them literal or nonliteral) motivated by compliance to that practice.

²⁵ At least to the same degree that languages are defeasible by recurrent violation. But see n. 24.

cooperation. In other words: we maintain the communal-scale cooperative communication by being truthful *simpliciter*, no matter the means.²⁶

Further, the possibility of using language as a communal means of communication is enabled by mutual trust that the community follows the general act *simpliciter*. Hence, speakers who betray the communal practice of telling the truth *simpliciter* either literally or nonliterally are bound to erode the possibility of literal and nonliteral communal communication.²⁷ Insofar telling the practice truth *simpliciter* is communally maintained and this practice enables the correct use of language, speakers who deceive nonliterally will betray the same kind of communal trust that is required to instance the literal use of linguistic conventions effectively.

Therefore, insofar as their communicative interactions are supported by the presumption of cooperation, deceiving speakers—whether through conventional or pragmatically rich means—can unfairly exploit a practice maintained by the efforts of the community. Hence, literal and nonliteral deception through cooperative signaling is systematically on par with respect to Fair Play in the plethora of cases in which successful communication is enabled by the communal efforts of a linguistic community. Speakers who effectively deceive by exploiting the presumption of cooperation are freeriding the communal efforts and sacrifices that maintain it.

Note too that this kind of consideration suggests a distinctive pathway to illuminate the connection between Fair Play and communication. Communal ventures almost always motivate communicative coordination. Maintaining the authority of the state or managing limited shared resources, for instance, often requires sharing information about the world. Hence, insofar as a communal venture is subject to Fair Play, communicative interactions motivated by this venture will be subject to Fair Play as well. Speakers who present themselves as committed to a beneficial communal venture in order to communicate in bad faith freeride the sacrifices required to preserve said communal venture. However, it would be arbitrary to posit a literalist restriction on this kind of communicative coordination.

Finally, there may be cases in which the presumption of cooperation is defeated, and interlocutors are still able to effectively use a communally maintained language. Say Omara and Kevin are CFOs of competing tech companies and are talking about where the best

²⁶ Arguably, the presumption of cooperation is preserved from erosion by a mixture of inherent social cognition (Boyd et al., 2003; Henrich, 2016; Sperber & Mercier, 2012; Tomasello, 2000, 2010) and external social norms that buttress the practice of telling the truth (Bicchieri, 2006; Graham, 2020; Murray & Starr, 2021; Starr, 2018).

²⁷ One may object that the literal deception is still morally worse off because she betrays Fair Play twice: first she betrays the communal presumption of cooperation; then, she betrays the communally built conventions that give meaning to her uttered sentences. More must be said about this to avoid objectionably double-counting wrongs. The liar betrays the presumption of cooperation by way of betraying the communally maintained practice of being cooperatively truthful. Hence, counting this as two wrongs would mischaracterize the type of the wrong as a token of the wrong.

processors are made. Say too that their companies are mutually threatened by a third, bigger company. Hence, despite their otherwise conflicting interests, both acknowledge that they are better off sharing information about where the best processors are made. While perhaps the presumption of communal cooperation is defeated by their generally antagonistic relationship, their local motivations in this case allow them to coordinate in communication. This local alignment allows Kevin to communicate that the best processors are made in Iquitos literally (“The best processors are made in Iquitos factory”) and nonliterally (“There’s a fantastic factory in Iquitos”).

Assume that, despite this local alignment, Kevin is still a devious speaker; there is a fantastic factory in Iquitos, but they are not making processors as of now. The speaker-community literalist may be tempted to argue that while his nonliteral deception betrayed the local commitment to cooperate, his nonliteral communication betrayed this same local commitment and the linguistic practices that enabled him to issue the linguistic conventions literally. Note, however, that this will not offer a systematic normative criterion that distinguishes between literal communication and nonliteral communication *tout court*. Rather, it will distinguish between literal communication and a limited subset of nonliteral communication—that which is not enabled nor serves to maintain the communal presumption of cooperation.

Berstler makes a strong point. The ethics of communication must include a communal lens. However, I have noted that i) there are plenty of communicative conventions that are not communally maintained, and ii) there is a pervasive and systematic way in which nonliteral communication is enabled by and preserves communal compliance with the practice of telling the truth *simpliciter* (be it literal or not). Hence, *if* Fair Play indeed governs communication, it does not systematically adjudicate between literal and nonliteral communication.

4. Cooperativism

I have argued against the claim that speaker-audience and speaker-community moral norms systematically differentiate literal from nonliteral communication. Yet, my aim is to substantiate the intuition that there is still a normative difference between committal and noncommittal communication. That is, there is still a normatively significant difference between Miguel’s committal utterance (1: She bribed the police) and his noncommittal utterance (2: She solved it like a businessman).

In this section I offer an alternative to Literalism, appealing to the notion of cooperative signaling. Some utterances inherit the instrumental commitments that enable cooperative signals. Others do not. This requires a more subtle parsing of the pragmatic realm. Nonliteral communication can be committal or noncommittal. Importantly, the normative difference I

characterize here is at least instrumental or procedural, stemming from the second-personal nature of cooperative action. Whether it is a moral difference depends on wider considerations on the ethics of cooperation that I discuss in §5.

4.1 Required Cooperative Interpretation

The arguments in §3 give us reason to think there is a unified speech act-kind of issuing cooperative signals. Agents who are motivated to act mutually and responsively to achieve a shared goal can issue utterances that are contextualized as communicating information about the world.

If you and I are committed to cooking pasta together and you are looking for the salt, I can issue the cooperative signal that the salt is in the cupboard by pointing at it and by saying “I put seasonings in the left one”, “the left one”, or “the salt is in the left cupboard”. Importantly, this interpretation is *required* to make sense of these utterances as cooperative moves in conversation. Given my commitment to coordinating with you, I can issue signals that must be interpreted as sharing information about the world if they are to have any role in the cooperative structure of our interaction. This, I submit, is the normative core of committal communication.

More schematically, I suggest that we can capture committal communication as follows:

Committal Communication: By issuing an utterance *u*, a speaker commits to a content *p* iff

interpreting *u* as aimed to inform that *p* is *required* by the best cooperative explanation(s) of how *u* plays a means towards the interlocutors’ acknowledged shared goals.

As defined, committal communication abstracts from the specific mechanisms of content recovery and is therefore ecumenic with respect to several modeling frameworks.²⁸ It also remains agnostic with respect to whether committal communication corresponds with other terms in the literature—say, testimony, tellings, sayings, or assertions.

What matters for my purposes is that cooperation plays a fundamental role in the speaker’s practical reasoning in issuing a signal and her audience’s epistemic reasoning in interpreting that signal. Hence, as we have seen, whether a content is cooperatively required is orthogonal to the literal or nonliteral means by which it is instanced. Sometimes, the

²⁸ For instance, one may render cooperative inferences in terms of question-answer dynamics (e.g., Roberts, 2012, 2018; Simons, 2010; Simons et al., 2010; Stokke 2012), game-theoretic modeling of coordination games (e.g. Jaegher, 2006; Franke, 2013; Martínez & Godfrey-Smith, 2016; Rothschild, 2013), or in terms of maxim-guided inference to the best explanation (e.g. Geurts, 2010; Levinson, 2000; Sauerland, 2004).

conversational context and the goals of the conversation will require literal interpretations; some other times, they will require nonliteral interpretations. Let me issue three clarifications to sharpen this view.

First Clarification: what counts as the best explanation of an utterance is a function of practical coherence, simplicity, and stability of an intentional action as means towards a given goal.²⁹ This allows for two or more interpretations to serve an equally good cooperative explanation of an utterance's role in conversation. In that case, the content to which speakers commit to is required by the intersection of the contents that serve as the best cooperative interpretations of said utterance.

Second clarification: the understanding of cooperation here is permissive. All it takes for agents to cooperate is that they mutually and responsively interlock their actions towards achieving the same goal. Hence, cooperative communication can be guided by a non-communicative shared goal—say, cooking pasta—or by a purely communicative goal—say, answering the question ‘where is Laurie Anderson from?’. By the same token, whether agents cooperate is independent of the reasons they have to do so. You may agree to cook with me because you are hungry or because you want to help a friend. However, you may also be motivated by the threat of what I would do if you did not. Likewise, committal communication can occur even in generally antagonistic interactions. Dueling noblemen may want to kill each other, and still be motivated to cooperate to find a place to duel. Nonideal communication, we will see in §5, has a vital role in illuminating the ethics of communicative commitments.

Third clarification: acknowledged shared goals need not be personally endorsed by interlocutors. A speaker may present herself as a cooperative interlocutor but fail to act accordingly. Most relevantly for us, speakers can present themselves as cooperative in bad faith in order to deceive their audience. I may present myself as committed to cooking pasta with you and appeal to that commitment to issue a cooperative signal communicating that the salt is in the cupboard. All this can be true, even if in fact all I wanted was to scare you with the fake spider I put there this morning.³⁰ Similarly, committal communication may also ensue when there is doubt that the speaker is merely acting *as if* she is committed to the goals that structure the conversation. In those cases, speakers can commit to contents that are required as cooperative interpretations with respect to said goals.

²⁹ See Bratman (1987), Bratman et al. (1988) Elster (2008), Harman, (1976, 1965) for considerations referring to action in general. See also Bratman (1992; 2014) Cohen & Levesque (1988; 1991), Litman & Allen (2003) for considerations on how coherence, stability, and simplicity apply to cooperative action.

³⁰ Thus, required contents can be communicated accidentally. The cooperative interpretation of my pointing to the cupboard is that the salt is there—no matter if my intention-in-action was to tell you that there may be spiders in the cabinet or that I really like that cabinet.

We can see now how the role of cooperative reasoning in enabling cooperative signals explains their distinctive normative import. To even be able to issue a cooperatively required content, speakers must present themselves as committed to a shared goal. In turn, they must issue an utterance such that interpreting it as aimed at communicating a given content is the best explanation of role as means towards said shared goal. Therefore, failing to provide the information one presents as providing amounts to failing to advance the very shared one has committed. It is, let us say, a constitutive failure; for it is a failure to advance the very same aims that enabled the speaker to issue the cooperative signal in the first place. If there is no salt in the cupboard, I have acted against our goal of cooking pasta together.

Hence, cooperatively required contents *inherit* the cooperative commitments that enable speakers to issue them. In both (3) and (4) Andrea has presented a signal as issuing information that will serve as means to get to the movies; hence, her commitment to going to the movies extends to her commitment to providing the means necessary to get there.

In turn, intentionally issuing a defective cooperative signal amounts to intentionally betraying the cooperative commitment that enabled said signal. Andrea's deception amounts to intentionally betraying her acknowledged commitment to go to the movies with Felipe. She presents herself as willing and able to act mutually and responsively towards their shared goal and, in virtue of that commitment, she exploits Felipe's trust that she will do her part in coordination in order to advance her own, incompatible, goals. To be effective, deceptive cooperative signals require the speaker's calling for the trust of her audience in bad faith.

I postpone discussion of the necessary and sufficient conditions of deceptive communication. Paradigmatically, a speaker issues a deceptive statement when she knows it is false and she has the intention that her audience forms a corresponding false belief.³¹ Clearly, this amounts to a cooperative failure. Yet, arguably, actually cooperative actions at least preclude believing that one will not do what one intends to do. Therefore, a statement is deceptive even when the speaker believes that it is false but turns out to be true. If Andrea believed that there was no gas at the station, but there actually was, she still acted against the commitment that enabled her cooperative signal.³²

³¹ Can speakers commit to contents publicly known to be false? Since—at least intuitively—this kind of deception is a weaker failure than fully committal deception, I bracket this issue in what follows. Nonetheless, consider a tentative negative answer. It is plausible that to even begin considering whether an act is cooperative or not it must be possible that it can serve as a means to achieve a shared goal. And since cooperatively required interpretations presume that communicating truth of a content serves as a means towards a shared goal, the fact that the cooperating agents know that it is false entails that clear to both interlocutors communicating its truth can be a means towards a shared goal.

³² This suggests why irresponsible communication amounts to a cooperative failure as well. Say that you and I are preparing a party and it's clear that my role will be to invite Zeyad. If my way to do so is dialing a random 10-digit number hoping that it will be his number, I am failing to be cooperative—even if by chance I dial his number. Similarly, stating content one is merely guessing (or being a covert bullshitter) is a cooperative failure.

What matters for our purposes, is that the facts in virtue of which speakers communicate a cooperatively required content p generate cooperative commitments towards p . In turn, intentionally issuing defective cooperatively required contents amounts to an intentional constitutive failure to comply with the cooperative commitments that make it possible to issue said content.

4.2 *Merely Allowed Cooperative Interpretation*

Recall our initial case. In casual conversation, Miguel communicates that Ana bribed the police, by issuing either (1: She bribed the police) or (2: She solved it like a businesswoman). The intuition to vindicate is that if it turns out that Miguel is deceiving his audience and Ana did not bribe the police, it seems that his uttering (2) can genuinely evade the criticisms that are entirely appropriate to his uttering (1). What explains this normative difference?

In short, Miguel and his audience know that (2) is compatible with several cooperative interpretations, given the vague goals that structure their office talk. And while these possible cooperative interpretations include that Ana bribed the police, they also include other innocuous interpretations—such as that she negotiated in noncriminal ways to be released. Therefore, the cooperative commitment that guides the conversation is not inherited by either of these specific interpretations.

We can capture noncommittal communication by appealing to cooperation as well:

Noncommittal Communication: By issuing an utterance u , a speaker communicates a content p without committing to it iff

- a) interpreting u as informing that p follows from an intentional interpretation of u that is *allowed* by the best cooperative explanation(s) of how u plays a means towards the interlocutors' acknowledged shared goals, but
- b) p is *not required* by said cooperative explanation(s).

Noncommittal communication is determined by how utterances fall short of playing a fully cooperative role in conversation. To be sure, Miguel must at least present himself as sharing some goals with his audience. Agents with no shared goals have no reason to even begin a conversation. If Miguel's utterance is to be interpreted as noncommittal communication and not a total communicative failure, it must make at least some sense that he intended to issue it as part of the structure of the conversation. Consequently, that Ana bribed the police must be allowed as a practically plausible interpretation that advances the goals of the conversation.

Nevertheless, this same content is *not required* by the goals of the conversation. That Ana bribed the police will not be the best explanation of Miguel's utterance as a cooperative move in conversation. Rather, as a communicative move aimed at advancing the goals that

structure the conversation, it stands on equal interpretive grounds to other, more innocuous interpretations. And, since cooperative interpretations are those required by the best explanations of an utterance conversational role as directed to achieve a shared goal, the required cooperative interpretation of (2) will be only what is required by *all* of these possible interpretations. This includes the fact that Ana did something businesswoman-like that led to her release, but not necessarily that it involved a bribe—perhaps she convinced the police of the costs of retaining her or convinced them that she was innocent given her impeccable criminal record.

The difference between committal and noncommittal communication can be seen as a function of whether a given interpretation has been sufficiently focused as a cooperative conversational move by the relationship between an utterance and the goals of the conversation. Committal contents are focused by the fact that they must be interpreted as such to make sense of an utterance as a cooperative move towards the shared goals that structure the conversation. Noncommittal contents are not so focused. Unlike Miguel's utterance (1)—and Andrea's utterances (3) and (4)—his utterance (2) exploits the interpretive space created by its vague instrumental relationship with the purposes of the conversation. Noncommittal communication exploits the instrumentally ambiguous role of a given utterance in advancing the goals of the conversation. Hence, while committal communication precludes cooperative ambiguity, noncommittal communication requires it.

Crucially, noncommittal communication still follows from intentional interpretations. In some cases, this means that audiences can interpret the speaker's intentions as they go beyond her acknowledged cooperative commitments. Given the vague goals that structure office talk and the public fact that Miguel has little motivation to cooperate his way into accusing his coworker of a crime, it will be clear that (2) falls short of advancing the cooperative commitments Miguel has with his audience. Nevertheless, Miguel's audience may still infer that he intends to insinuate that Ana bribed the police. However, the inference that leads to this conclusion will not hinge on the cooperative commitments that Miguel has acknowledged. Rather, such an inference may rely on the facts that Miguel dislikes Ana or that Ana has a pattern of bribing authority figures.

In other—and to my ear, more common—cases, speakers can instance noncommittal communication by merely gesturing to a vague range of plausible allowed contents without it being clear which one specifically they mean. Sometimes, that is, there is no specific intended content that can be deciphered but an ill-defined set of contents that could be intended by the speaker. Say that Catherine and Felicia are deciding whether to end their date. Say too that Felicia utters “there are many things we could do” in a flirting tone. In that case, Catherine can understand Felicia's intention to issue an utterance that can be interpreted in a variety of

ways, including that they should continue their date somewhere else or go to her place, but it is also compatible with her suggestion to just go home.³³

Note that, along Gricean tenets, the present view gives cooperation an essential role in determining what utterances count as committal or noncommittal communication. However, against classic Gricean views, noncommittal communication shows that not every speaker intention stands equally. As Elizabeth Camp (2019) argues, noncommittal insinuations put pressure on classic Gricean pictures because they show that recognition of communicative intentions can have a noncooperative role in conversation.³⁴ Insinuating speakers can tailor their utterances such that their audiences can recognize their intentions to insinuate a content and yet be unable to hold accountable speakers for having committed to that content. If I am right, this commitment profile owes to the fact that—unlike committal communication—the individual intentions that guide noncommittal communication are not an essential part of the cooperative structure that guides the conversation.

Hopefully, it is clear how Cooperativism explains the normative distinction between committal and noncommittal communication. Noncommittal contents are not required as cooperative interpretations of an utterance as a means to the interlocutors' shared goal. In fact, several instances of noncommittal communication are successful because audiences can recognize that the communicated content does not fulfill this cooperative role.

Hence, in our terms, noncommittal contents do not inherit the cooperative commitments to the shared goals that structure the conversation. Miguel's commitment to carry on with his casual conversation does not commit him to the content that Ana bribed the police when he utters (2)—while it does when he utters (1). Consequently, deceptive noncommittal contents, unlike deceptive committal contents, do not constitutively contravene the cooperative commitments that structure the conversation with her audience. Therefore, intentionally issuing a false noncommittal content does not amount to the intentional betrayal of the cooperative commitments that structure the conversation.

Of course, if his noncommittal insinuation is deceptive, Miguel has arguably breached other norms. Communicating false contents through deceptive insinuations does not absolve speakers of any kind of accountability. The fact that he intentionally (even if nont cooperatively) communicated false content may constitute a breach of norms that involve harm or respect of other's informed autonomy—for instance, norms governing what Michaelson and Stokke (2022) call positional advantages. Nonetheless, these further norms

³³ In a sense, then, noncommittal communication can have a cooperative role of somehow advancing the shared goals of the interlocutors. This role, however, is not that of communicating a cooperatively required proposition.

³⁴ Note that 'noncooperative' does not amount to 'uncooperative'. A cooperative intention is an intention that is subordinated to a higher-order intention to achieve a shared goal. A non-cooperative intention is not so subordinated. An uncooperative intention goes against a shared goal.

will not explain the normative difference between committal and noncommittal communication, as they apply to both.

Admittedly, the threshold between required and merely allowed contents is sometimes vague and murky. In some cases, it may not be clear whether a speaker has committed to a given content. Nonetheless, the distinction and its explanatory potential is clear in paradigmatic cases.³⁵

Further, note that different motivational alignments may render what was a committal content in one context into a noncommittal one in another context.³⁶ Say that Andrea and David were not going to the movies, but in a competitive scavenger hunt only one of them can win. Given each agent's motivations, there are good reasons to think that Andrea is not interested in getting Felipe some gas. Arguably, this may block her implicature in (2) from being cooperatively required and thus from committing her to the fact that Felipe can get gas at the station down the road. Yet, her utterance in (1) may still commit her to this content even in this contentious context. After all, given the public conventional meaning of her utterance, there are no other plausible cooperative interpretations of her words.

Hence, in some cases literal communication can indeed play the role of ensuring safe cooperative interpretations. This is why in contentious conversations—as in conversations with the police or in competitive business negotiations—audiences are warranted in pressuring speakers to be as literally explicit as possible (“but do they sell gas in that station?”). Nevertheless, extremely minimal cooperation in conversation may even block the contextual supplementation that is intuitively involved in interpreting literal contents. In some cases, it may be appropriate for audiences to press uncooperative speakers into making explicit their use of relative adjectives (“how tall are you saying Yuri is”), relational expressions (“Paulo is ready for what exactly?”), pronouns (“Who do you mean by ‘he’?”), and even ambiguous terms (“Do you mean Geneva, Illinois, or Geneva, New York?”).

Therefore, weak or contentious cooperation does not yield a strict normative distinction between literal and nonliteral communication. What matters is whether a given utterance warrants a cooperative interpretation, given the public resources available to the interlocutors. This gives us no reason to eschew the plethora of cases in which fully cooperative engagements warrant committal nonliteral communication. In fact, this is an extensional advantage of cooperativism; for it captures the ways variation in coordination contexts can shift what communicative strategies can yield communicative commitments. It should not be surprising that a conversation with a hostile police officer will allow different communicative moves than a conversation with an age-old friend.

³⁵ Note that Literalism is not necessarily better in this respect. As I have noted, there are plenty of contentious cases in which it is not clear whether communicated contents are literally encoded in a given uttered sentence.

³⁶ See Thomason (2003) for a brief discussion on the effects of contentious dynamics on communicated contents.

5. The Ethics of Communicative Commitment

Recall the questions we began with:

- A) How should we characterize the distinction between committal and noncommittal speech (if there is one at all)?
- B) Why or in virtue of what is this distinction normatively significant (if it is at all)?

I have answered (A) by appealing to the distinction between required and merely allowed contents. In response to (B), I argued that the facts in virtue of which speakers are able to issue cooperative signals generate commitments towards cooperatively required contents that are not carried by merely allowed contents. These commitments are instrumental or procedural, stemming from the requirements of joint action. Failure to honor these commitments amounts to a failure to do what one has committed to do with others. These are second-personal failures, as they involve a kind of action that can only occur between agents that recognize each other as agents—willing and able to see each other as co-authors of a distinctive kind of action (e.g., cooking *together*, going to the movies *together*.)

Yet, the question remains, does intentionally failing to do what one committed to do with others amount to a *moral* failure beyond the instrumental requirements of joint action? So far, Cooperativism is less ambitious than Literalism. It holds that the normative significance between required and merely allowed cooperative contents is instrumental and second-personal; yet it need not be moral.

Note, however, that Cooperativism yields a distinctive framework for theorizing about the ethics of communicative commitments. It shifts the ethics of communicative commitment from considerations about the proper or improper use of our words to considerations about the moral significance of joint agency. One may argue that commitments to act with others inherently generate moral reasons to comply with them. Therefore, there are moral (even if only *pro tanto*) reasons not to betray the cooperative commitments that enable cooperatively required contents. Call this position *internalism* concerning the moral significance of cooperation. On the other hand, one may adopt an *externalist* stance about the moral status of cooperative commitments. One may think that, since cooperation can be motivated in very different ways, whether cooperative commitments generate moral reasons to comply with them depends on factors external to said commitments.

I do not aim to adjudicate between these alternatives here. This will require important theorizing about the ethics of joint action. Rather, I conclude by discussing how the contention between internalism and externalism shows that cooperativism offers important tools to account for the individual and collective ethics of communication in ways that literalism is just too narrow to capture. In particular, I focus on how cooperativism illuminates

the role of higher-order social phenomena—such as social power, norms, and institutions—in thinking about the ethics of communicative commitments.

Arguably, the internalist may simply adapt the literalist's arguments to apply to cooperatively required contents instead of literal contents. As I have argued, successful cooperative action demands trust that others will act in ways that allow us to mutually and responsively achieve what we intend. One may think that this trust, and not the trust that interlocutors are properly using semantic conventions, generates the kind of moral considerations that the speaker-audience literalist appealed to. These considerations of autonomy and trust would then generate distinctive moral reasons not to deceive through cooperatively required contents (literal or not).³⁷

Further, the internalist may argue that the considerations of speaker-audience literalism may also extend to cooperatively required contents. Exploiting commitments to beneficial communal endeavors to communicate in bad-faith breaches fairness considerations. Arguably, one such communal endeavor is maintaining the presumption of communicative coordination. Hence the wide range of literal and non-literal communicative commitments that rely on this presumption will be subject to fair play considerations. But more generally, virtually any non-communicative communal endeavor that is governed by fair play will most plausibly motivate communicative coordination at a communal scale—rendering the exploitation of communicative commitments embedded in these endeavors unfair to the community that maintains them.

Most ambitiously, these considerations can be unified by noting that the distinctive abilities to act with others in ways that allow for local and large-scale communicative coordination constitute a fundamental aspect of human agency. On the one hand, we gain the abilities to regard others as agents through developing the abilities to coordinate with them.³⁸ On the other hand, we developed distinctively human large-scale coordination because we generated just and beneficial cultural, social, and political devices that ensure stable and beneficial coordination between individual agents.³⁹ Hence, insofar as they constitute an vital part of our humanity, commitments to act with others generate moral reasons to comply with

³⁷ For example, Shiffrin's (2016) arguments concerning linguistic deception. If cooperative signaling and not literally speaking is fundamental to glean each other's minds, her arguments concerning the moral status of deception would arguably extend to beyond literalism.

³⁸ The idea that humans have distinctive cognitive skills for joint action has been defended at least since Fichte (2000) (see Kosch 2021). See also Mead (2000) and Tomasello (2019).

³⁹ See Henrich (2016) for a thorough review of how human gene-culture coevolution ensured that large scale social phenomena that ensured fitness and coordination within human communities. The internalist may appeal to these mechanisms and the plausible assumption that at least some of them have a role in maintaining fair and just communal endeavors.

them. Agents who cannot act together in the ways that enable cooperative signals, cannot enact the kind of agency that distinguishes our humanity.

In fact, cooperativism yields a novel view of the commonalities between the harms entailed by epistemic and communicative injustices. In a plethora of cases, social norms, biases, and institutions prevent marginalized individuals from participating in the kinds of second-personal coordination that foster full-blooded human agency. In terms of communication, these social factors hinder people from effectively speaking their mind, and even from generating communicative commitments altogether. Sometimes, this happens because the speakers' public authority to be recognized as willing and able to coordinate in communication is diminished in virtue of their social standing, resulting in testimonial injustices.⁴⁰ Other times, the social standing of speakers renders them simply unable to coordinate with their audiences in the ways they want to, leading to silenced speech.⁴¹

However, the internalist may be overly optimistic. Human coordination is often and systematically non-ideal. All it takes to cooperate is that two or more agents mutually and responsively interlock their actions to achieve the same mutually acknowledged aim. Such an aim, however, may be motivated in several ways. You and I may be coordinating to make pasta because we are friends. Yet, we may also be motivated by the fact that it is our job, or even because you are afraid of what I would do if you did not help. By the same token, joint action can be aimed towards nefarious aims. Destructive military operations, land embezzlement, and misinformation campaigns are more often than not carried out through coordination.

Moreover, consider the fact human communities generate higher-order social mechanisms to ensure the coordination of agents that would not be motivated to coordinate otherwise. Social power and reputation costs often motivate speakers to comply with the demands of powerful agents.⁴² Communities develop patterns of informal social sanctions, expectations, and norms that motivate coordination from agents that would not coordinate otherwise.⁴³ And, institutional procedures enable and generate distinctive kinds of social acts that would not be available without social compliance—just as the act of applying for a mortgage, baptizing your child, or testifying in court.⁴⁴

⁴⁰ See Fricker (2009), Jugov (2019). See Crawford (forthcoming) for an account of testimonial injustice as a failure of recognition. See Hawley (2011) for the practical dimension of epistemic injustice in terms of know-how.

⁴¹ See Langton (1993), Maitra (2009), Dotson (2011), McGowan (2019). See Wieland (2007) and Unnsteinsson (2019) for discussion on the role of convention in silencing.

⁴² See Bierstedt (1950) and French & Raven (1959) for classic accounts of social power as a particular kind of influence in sociology and Friedkin (2010) for a formal development of French's framework. See Fiske & Berdahl (2007) for an account of power as control over another's valued outcomes.

⁴³ See Bicchieri (2006) and Gavrillets & Richerson (2017) for accounts of social norms as involving expectations of social sanctions. See Brennan et al. (2016) for an unconditional account of social norms.

⁴⁴ See Searle (2017), Guala (2016).

These higher-order social mechanisms, however, often and systematically harm the agents they govern and thus preserve arrangements that enact and maintain social injustice. It is easy to imagine how agents may use their power to motivate communicative cooperation from those in subordinated positions. The fact that Donald knows that Jane understands the very real consequences of disobeying him, given the power he wields over her, may warrant his trust in Jane's testimony that the documents have been destroyed. Similarly, Elinor Ochs and her colleagues (1993) report how social norms of speech in certain communities shift the burden of proof and the sanctions for defective assertions onto those who are lower in the social hierarchy. Social norms motivate and constraint communicative coordination that map into and preserve the social hierarchies of a community. And consider too John Rickford and Sharese King's (2016) evidence that Black communities in the US are often forced to be part of legal systems that linguistically disadvantage them in court. Black communities are forced to abandon vernacular speech to even be able to participate in the institutionalized acts that constitute the country's legal institutions.

Joint action may be all too human after all.⁴⁵ It may distinguish human agency; but precisely for this reason it also enables distinctively human ways of harming each other. In fact, we have just seen, cooperativism not only accounts for injustices that occur when agents are excluded from second-personal ways of acting together; it also accounts for a distinctive—and often neglected—kind of injustice that occurs when agents are unduly brought *into* these same kind of second-personal relationships. Hence, the externalist may argue, even if the commitments that enable cooperative signals generate second-personal instrumental reasons to comply with these commitments, the fact that these commitments are often and systematically hijacked by unjust social phenomena may give us pause in thinking that there are moral reasons to comply with these commitments. If so, whether deception through cooperative communication amounts to a moral failure, will come down to external considerations concerning the contexts that motivate our communicative interactions.

Whether the internalist/externalist contention lands, cooperativism yields fertile grounds to understand the moral stakes of communicative commitments in ways that extend the limits of the proper or improper use of the conventional tools we develop to communicate. By shifting the normative focus to the conditions in which we act with each other in communication and the reasons interlocutors may have to do so, cooperativism provides a distinctive link between the ethical import of our communicative practices and the inescapable sociality of human agency.

⁴⁵ Cf. Manne (2016; 2017).

References

- Adler, J. E. (1997). Lying, Deceiving, or Falsely Implicating. *The Journal of Philosophy*, 94(9), 435–452. <https://doi.org/10.2307/2564617>
- Adler, J. E. (2018). Lying and Misleading: A Moral Difference. In Lying (pp. 301–318). Oxford University Press. <https://oxford.universitypressscholarship.com/view/10.1093/oso/9780198743965.001.0001/oso-9780198743965-chapter-16>
- Adler, J. E. (2013). *Epistemological Problems of Testimony*. Oxford University Press.
- Arneson, R. J. (1982). The Principle of Fairness and Free-Rider Problems. *Ethics*, 92(4), 616–633.
- Augustine. (2002a). Lying. In R. J. Deferrari (Ed.), & M. S. Muldowney, M. E. DeFerrari, M. F. McDonald, L. Meagher, & C. Eagan (Trans.), *Treatises on Various Subjects* (1st pbk. reprint). Catholic University of America Press.
- Augustine. (2002b). *Treatises on Various Subjects* (R. J. Deferrari, Ed.; M. S. Muldowney, M. E. DeFerrari, M. F. McDonald, L. Meagher, & C. Eagan, Trans.; 1st pbk. reprint). Catholic University of America Press.
- Bach, K. (1997). The Semantics-Pragmatics Distinction: What It Is and Why It Matters. In E. Rolf (Ed.), *Pragmatik: Implikaturen und Sprechakte* (pp. 33–50). VS Verlag für Sozialwissenschaften. https://doi.org/10.1007/978-3-663-11116-0_3
- Bach, K. (2005). Context Ex Machina. In Z. G. Szabo (Ed.), *Semantics Versus Pragmatics* (pp. 15–44). Oxford University Press.
- Bach, K. (2006). The Excluded Middle: Semantic Minimalism Without Minimal Propositions. *Philosophy and Phenomenological Research*, 73(2), 435–442. <https://doi.org/10.1111/j.1933-1592.2006.tb00626.x>
- Berstler, S. (2019). What's the Good of Language? On the Moral Distinction Between Lying and Misleading. *Ethics*, 130(1), 5–31. <https://doi.org/10.1086/704341>
- Bicchieri, C. (2006). *The Grammar of Society: The Nature and Dynamics of Social Norms*. Cambridge University Press.
- Bierstedt, R. (1950). An Analysis of Social Power. *American Sociological Review*, 15(6), 730–738. <https://doi.org/10.2307/2086605>
- Borg, E. (2007). Minimalism versus Contextualism in Semantics. In G. Preyer & G. Peter (Eds.), *Context-Sensitivity and Semantic Minimalism: New Essays on Semantics and Pragmatics*. Oxford University Press.
- Borg, E. (2010). Minimalism and the Content of the Lexicon. In E. Rast & L. C. Baptista (Eds.), *Meaning and Context* (pp. 51–77). Peter Lang.
- Borg, E. (2019). Explanatory Roles for Minimal Content. *Noûs*, 53(3), 513–539.
- Boyd, R., Gintis, H., Bowles, S., & Richerson, P. (2003). Evolution of Altruistic Punishment. *Proceedings of the National Academy of Sciences*, 100, 3531–3535. <https://doi.org/10.1073/pnas.0630443100>
- Bratman, M. (1987). *Intention, Plans, and Practical Reason* (Issue 2, pp. 198–199). Cambridge, MA: Harvard University Press.

- Bratman, M. E. (1992). Shared Cooperative Activity. *The Philosophical Review*, 101(2), 327–341. JSTOR. <https://doi.org/10.2307/2185537>
- Bratman, M. E. (2014). Shared Agency: A Planning Theory of Acting Together. In *Shared Agency*. Oxford University Press. <https://www.oxfordscholarship.com/view/10.1093/acprof:oso/9780199897933.001.0001/acprof-9780199897933>
- Bratman, M. E., Israel, D. J., & Pollack, M. E. (1988). Plans and resource-bounded practical reasoning. *Computational Intelligence*, 4(3), 349–355. <https://doi.org/10.1111/j.1467-8640.1988.tb00284.x>
- Brennan, G., Eriksson, L., Goodin, R. E., & Southwood, N. (2016). *Explaining Norms*. Oxford University Press.
- Burge, T. (1997). Interlocution, Perception, and Memory. *Philosophical Studies*, 86(1), 21–47. <https://doi.org/10.1023/A:1004261628340>
- Camp, E. (2019). Insinuation, Common Ground, and the Conversational Record. In D. Fogal & D. W. Harris (Eds.), *New Work On Speech Acts*. Oxford University Press. <https://www.oxfordscholarship.com/view/10.1093/oso/9780198738831.001.0001/oso-9780198738831-chapter-2>
- Cappelen, H., & Lepore, E. (2005). *Insensitive Semantics: A Defense of Semantic Minimalism and Speech Act Pluralism* (1st edition). Wiley-Blackwell.
- Carston, R. (1999). The Semantics/Pragmatics Distinction: A View From Relevance Theory. In K. Turner (Ed.), *The Semantics/Pragmatics Interface From Different Points of View* (p. 85-125). Elsevier.
- Chisholm, R. M., & Feehan, T. D. (1977). The Intent to Deceive. *The Journal of Philosophy*, 74(3), 143–159. <https://doi.org/10.2307/2025605>
- Christiansen, M. H., & Chater, N. (2022). *The Language Game: How Improvisation Created Language and Changed the World*. Basic Books.
- Cohen, P. R., & Levesque, H. J. (1988). Rational Interaction as the Basis for Communication. *Pattern Recognition - PR*, 39.
- Cohen, P. R., & Levesque, H. J. (1991). Teamwork. *Noûs*, 25(4), 487–512. <https://doi.org/10.2307/2216075>
- de Jaegher, K. (2006). Game-Theoretic Grounding. In A. Benz, G. Jäger, & R. van Rooij (Eds.), *Game Theory and Pragmatics* (pp. 220–247). Palgrave Macmillan UK. https://doi.org/10.1057/9780230285897_8
- Crawford, L. (forthcoming). Testimonial Injustice and Mutual Recognition. *Ergo: An Open Access Journal of Philosophy*.
- Dotson, K. (2011). Tracking Epistemic Violence, Tracking Practices of Silencing. *Hypatia*, 26(2), 236–257. <https://doi.org/10.1111/j.1527-2001.2011.01177.x>
- Elster, J. (2008). *Reason and Rationality*. In Reason and Rationality. Princeton University Press. <https://doi.org/10.1515/9781400833177>
- Franke, M. (2013). Game Theoretic Pragmatics. *Philosophy Compass*, 8(3), 269–284. <https://doi.org/10.1111/phc3.12015>
- Fricke, E. (2012). Stating and Insinuating. *Aristotelian Society Supplementary Volume*, 86(1), 61–94. <https://doi.org/10.1111/j.1467-8349.2012.00208.x>
- Fricke, M. (2009). *Epistemic Injustice: Power and the Ethics of Knowing* (1 edition). Oxford University Press, USA.

- Friedkin, N. E. (1986). A formal theory of social power. *The Journal of Mathematical Sociology*, 12(2), 103–126. <https://doi.org/10.1080/0022250X.1986.9990008>
- French Jr., J. R. P., & Raven, B. (1959). The bases of social power. In *Studies in social power* (pp. 150–167). University of Michigan.
- Fichte, J. G. (2000). *Foundations of Natural Right: According to the principles of the Wissenschaftslehre* (F. Neuhouser & M. Baur, Trans.). Cambridge University Press.
- Fiske, S. T., & Berdahl, J. (2007). Social power. In *Social psychology: Handbook of basic principles*, 2nd ed (pp. 678–692). The Guilford Press.
- Gavrilets, S., & Richerson, P. J. (2017). Collective action and the evolution of social norm internalization. *Proceedings of the National Academy of Sciences of the United States of America*, 114(23), 6068–6073. <https://doi.org/10.1073/pnas.1703857114>
- Geurts, B. (2010). *Quantity implicatures* (1. publ). Cambridge Univ. Press.
- Geurts, B. (2018). Convention and common ground. *Mind & Language*, 33(2), 115–129. <https://doi.org/10.1111/mila.12171>
- Goldberg, S. C. (2007). *Anti-Individualism: Mind and Language, Knowledge and Justification*. Cambridge University Press. <https://doi.org/10.1017/CBO9780511487521>
- Goldberg, S. C. (2008). The Epistemic Utility of What is Said. In R. Stainton & C. Viger (Eds.), *Compositionality, Context, and Semantic Values: Essays in Honor of Ernie Lepore*. Springer.
- Gormally, M. C. (1991). The Ethical Root of Language. In P. Geach & J. Holowka (Eds.), *Logic and Ethics* (pp. 49–70). Springer Netherlands. https://doi.org/10.1007/978-94-011-3352-4_4
- Graham, P. J. (2020). The Function of Assertion and Social Norms. In S. Goldberg (Ed.), *The Oxford Handbook of Assertion* (pp. 725–748). Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780190675233.013.30>
- Green, S. (2001). Lying, Misleading, and Falsely Denying: How Moral Concepts Inform the Law of Perjury, Fraud, and False Statements. *Hastings Law Journal*, 53(1), 157.
- Grice, H. P. (1967). Logic and Conversation. In P. Grice (Ed.), *Studies in the Way of Words* (pp. 41–58). Harvard University Press.
- Grice, H. P. (1968). Utterer's Meaning, Sentence-Meaning, and Word-Meaning. *Foundations of Language*, 4(3), 225–242.
- Guala, F. (2016). *Understanding Institutions: The Science and Philosophy of Living Together* (1st edition). Princeton University Press.
- Harman, G. H. (1976). Practical Reasoning. *The Review of Metaphysics*, 29(3), 431–463. JSTOR.
- Harman, G. H. (1965). The Inference to the Best Explanation. *Philosophical Review*, 74(1), 88–95. <https://doi.org/10>
- Hart, H. L. A. (1955). Are There Any Natural Rights? *The Philosophical Review*, 64(2), 175–191. <https://doi.org/10.2307/2182586>
- Hawley, K. (2011). Knowing How and Epistemic Injustice. In J. Bengson & M. A. Moffett (Eds.), *Knowing How: Essays on Knowledge, Mind, and Action* (Issue 1, pp. 283–299). Oxford University Press.

- Hawthorne, J. (2012). Some Comments on Fricker's 'Stating and Insinuating.' *Aristotelian Society Supplementary Volume*, 86(1), 95–108. <https://doi.org/10.1111/j.1467-8349.2012.00209.x>
- Henrich, J. (2016). *The Secret of Our Success: How culture is driving human evolution, domesticating our species and making us smarter*. Princeton University Press.
- Jugov, T., & Ypi, L. (2019). Structural Injustice, Epistemic Opacity, and the Responsibilities of the Oppressed. *Journal of Social Philosophy*, 50(1), 7–27. <https://doi.org/10.1111/josp.12268>
- Kant, I. (1996). On a Supposed Right to Lie from Philanthropy. In M. J. Gregor (Trans.), *Practical Philosophy* (pp. 605–616). Cambridge University Press. <https://doi.org/10.1017/CBO9780511813306.014>
- Korta, K., & Perry, J. (2006). Varieties of Minimalist Semantics. *Philosophy and Phenomenological Research*, 73(2), 451–459. <https://doi.org/10.1111/j.1933-1592.2006.tb00628.x>
- Kosch, M. (2021). Fichte on Summons and Self-Consciousness. *Mind*, 130(517), 215–249. <https://doi.org/10.1093/mind/fzaa001>
- Langton, R. (1993). Speech Acts and Unspeakable Acts. *Philosophy & Public Affairs*, 22(4), 293–330.
- Lepore, E., & Stone, M. (2015a). David Lewis on Convention. In *A Companion to David Lewis* (pp. 313–327). John Wiley & Sons, Ltd. <https://doi.org/10.1002/9781118398593.ch20>
- Lepore, E., & Stone, M. (2015b). *Imagination and Convention: Distinguishing Grammar and Inference in Language*. Oxford University Press.
- Levinson, S. C. (2000). *Presumptive Meanings: The Theory of Generalized Conversational Implicature*. MIT Press.
- Lewis, D. (1975). Languages and Language. In K. Gunderson (Ed.), *Minnesota Studies in the Philosophy of Science* (pp. 3–35). University of Minnesota Press.
- Lewis, D. (1980). *Convention*. Blackwell.
- Litman, D. J., & Allen, J. F. (2003). Discourse Processing and Commonsense Plans. In P. R. Cohen, J. Morgan, & M. E. Pollack (Eds.), *Intentions in Communication*. The MIT Press. <https://doi.org/10.7551/mitpress/3839.003.0019>
- MacFarlane, J. (2011). What Is Assertion. In J. Brown & H. Cappelen (Eds.), *Assertion*. Oxford University Press.
- MacIntyre, A. (2006). Truthfulness and lies: What can we learn from Kant? In *Ethics and Politics: Selected Essays* (Vol. 2, pp. 122–142). Cambridge University Press. <https://doi.org/10.1017/CBO9780511606670.008>
- Mahon, J. E. (2003). Kant on Lies, Candour and Reticence. *Kantian Review*, 7, 102–133. <https://doi.org/10.1017/S1369415400001758>
- Marmor, A. (2009). *Social Conventions: From Language to Law*. Princeton University Press.
- Martínez, M., & Godfrey-Smith, P. (2016). Common Interest and Signaling Games: A Dynamic Analysis. *Philosophy of Science*, 83(3), 371–392. <https://doi.org/10.1086/685743>
- Maitra, I. (2009). Silencing Speech. *Canadian Journal of Philosophy*, 39(2), 309–338.

- Manne, K. (2016). Humanism: A Critique. *Social Theory and Practice*, 42(2), 389–415. <https://doi.org/10.5840/soctheorpract201642221>
- Manne, K. (2017). *Down Girl: The Logic of Misogyny*. Oxford University Press.
- Mazzarella, D., Reinecke, R., Noveck, I., & Mercier, H. (2018). Saying, presupposing and implicating: How pragmatics modulates commitment. *Journal of Pragmatics*, 133, 15–27. <https://doi.org/10.1016/j.pragma.2018.05.009>
- McGowan, M. K. (2017). On Multiple Types of Silencing. In M. Mikkola (Ed.), *Beyond Speech: Pornography and Analytic Feminist Philosophy* (p. 0). Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780190257910.003.0003>
- Mead, G. H. (2000). *Mind, self, and society: From the standpoint of a social behaviorist* (C. W. Morris, Ed.). Univ. of Chicago Press.
- Michaelson, E., & Stokke, A. (2021). Lying, Deception, and Epistemic Advantage. In *The Routledge Handbook of Social and Political Philosophy of Language*. Routledge.
- Michaelson, E. (2013). *This and That: A Theory of Reference for Names, Demonstratives, and Things in Between* [PhD Thesis]. UCLA.
- Millikan, R. G. (1987). *Language, Thought, and Other Biological Categories: New Foundations for Realism* (Reprint edition). MIT Press.
- Murray, S. E., & Starr, W. B. (2021). The Structure of Communicative Acts. *Linguistics and Philosophy*, 1–50. <https://doi.org/10.1007/s10988-019-09289-0>
- Nozick, R. (2013). *Anarchy, State, and Utopia* (Reprint edition). Basic Books.
- Ochs, E., Duranti, A., & Goodwin, C. (1993). Indexing gender. *Sex and Gender Hierarchies*, 146–169.
- Owens, D. (2017). Human Testimony. In *Normativity and Control*. Oxford University Press. <https://doi.org/10.1093/oso/9780198713234.003.0011>
- Owens, D. (2022). *Bound by Convention: Obligation and Social Rules*. Oxford University Press.
- Peet, A. (2015). Testimony, Pragmatics, and Plausible Deniability. *Episteme*, 12(1), 29–51. <https://doi.org/10.1017/epi.2014.31>
- Peet, A. (2016). Testimony and the epistemic uncertainty of interpretation. *Philosophical Studies*, 173(2), 395–416. <https://doi.org/10.1007/s11098-015-0498-x>
- Pepp, J. (2019). Assertion, Lying, and Untruthfully Implicating. In S. C. Goldberg (Ed.), *The Oxford Handbook of Assertion*. Oxford University Press.
- Perry, J. (2007). Situating Semantics: A Response. In M. O. C. Washington (Ed.), *Situating Semantics: Essays on the Philosophy of John Perry* (pp. 507–575).
- Rawls, J. (1999). *A Theory of Justice* (2nd edition). Belknap Press: An Imprint of Harvard University Press.
- Recanati, F. (2003). *Literal Meaning*. Cambridge University Press.

- Reins, L. M., & Wiegmann, A. (2021). Is Lying Bound to Commitment? Empirically Investigating Deceptive Presuppositions, Implicatures, and Actions. *Cognitive Science*, 45(2), e12936. <https://doi.org/10.1111/cogs.12936>
- Rickford, J. R., & King, S. (2016). Language and linguistics on trial: Hearing Rachel Jeantel (and other vernacular speakers) in the courtroom and beyond. *Language*, 92(4), 948–988. <https://doi.org/10.1353/lan.2016.0078>
- Roberts, C. (2012). Information structure in discourse: Towards an integrated formal theory of pragmatics. *Semantics and Pragmatics*, 5. <https://doi.org/10.3765/sp.5.6>
- Roberts, C. (2018). Speech Acts in Discourse Context. In *New Work on Speech Acts*. Oxford University Press. <https://www.oxfordscholarship.com/view/10.1093/oso/9780198738831.001.0001/oso-9780198738831-chapter-12>
- Rothschild, D. (2013). Game Theory and Scalar Implicatures. *Philosophical Perspectives*, 27(1), 438–478. <https://doi.org/10.1111/phpe.12024>
- Sauerland, U. (2004). Scalar Implicatures in Complex Sentences. *Linguistics and Philosophy*, 27(3), 367–391. <https://doi.org/10.1023/B:LING.0000023378.71748.db>
- Saul, J. M. (2002). Speaker Meaning, What is Said, and What is Implicated. *Noûs*, 36(2), 228–248. <https://doi.org/10.1111/1468-0068.00369>
- Saul, J. M. (2012). Lying, Misleading, and What is Said: An Exploration in Philosophy of Language and in Ethics. In *Lying, Misleading, and What is Said*. Oxford University Press. <https://www.oxfordscholarship.com/view/10.1093/acprof:oso/9780199603688.001.0001/acprof-9780199603688>
- Schelling, T. C. (1981). *The Strategy of Conflict: With a New Preface by the Author* (Reprint edition). Harvard University Press.
- Searle, J. R. (1970). *Speech Acts: An Essay in the Philosophy of Language*. Cambridge University Press.
- Shiffrin, S. V. (2016). *Speech Matters: On Lying, Morality, and the Law* (Reprint edition). Princeton University Press.
- Simmons, A. J. (1979). The Principle of Fair Play. *Philosophy & Public Affairs*, 8(4), 307–337.
- Simons, M. (2010). A Gricean View on Intrusive Implicatures. In K. Petrus (Ed.), *Meaning and Analysis: New Essays on Grice* (pp. 138–169). Palgrave Macmillan UK. https://doi.org/10.1057/9780230282117_6
- Simons, M., Tonhauser, J., Beaver, D., & Roberts, C. (2010). What projects and why. *Semantics and Linguistic Theory*, 20(0), 309–327. <https://doi.org/10.3765/salt.v20i0.2584>
- Sperber, D., & Mercier, H. (2012). Reasoning as a Social Competence. In H. Landemore & J. Elster (Eds.), *Collective Wisdom: Principles and Mechanisms* (pp. 368–392). Cambridge University Press. <https://doi.org/10.1017/CBO9780511846427.016>
- Stanley, J., & Szabó, Z. G. (2000). On Quantifier Domain Restriction. *Mind & Language*, 15(2–3), 219–261. <https://doi.org/10.1111/1468-0017.00130>
- Starr, W. B. (2018). *Socializing Pragmatics* (Vol. 1). Oxford University Press. <https://doi.org/10.1093/oso/9780198791492.003.0007>
- Stokke, A. (2016). Lying and Misleading in Discourse. *The Philosophical Review*, 125(1), 83–134. <https://doi.org/10.1215/00318108-3321731>

- Stokke, A. (2019). Lies, Harm, And Practical Interests. *Philosophy and Phenomenological Research*, 98(2), 329–345. <https://doi.org/10.1111/phpr.12439>
- Stokke, A., & Schoubye, A. J. (2016). What is Said? *Noûs*, 50(4), 759–793. <https://doi.org/10.1111/nous.12133>
- Strudler, A. (2010). The Distinctive Wrong in Lying. *Ethical Theory and Moral Practice*, 13(2), 171–179.
- Thomason, R. (2003). Accommodation, Meaning, and Implicature: Interdisciplinary Foundations for Pragmatics. In P. R. Cohen, J. Morgan, & M. E. Pollack (Eds.), *Intentions in Communication*. (pp. 235–363). MIT Press.
- Timmermann, F., & Viebahn, E. (2021). To lie or to mislead? *Philosophical Studies*, 178(5), 1481–1501. <https://doi.org/10.1007/s11098-020-01492-1>
- Tomasello, M. (2000). First steps toward a usage-based theory of language acquisition. *Cognitive Linguistics*, 11(1–2), 61–82. <https://doi.org/10.1515/cogl.2001.012>
- Tomasello, M. (2019). *Becoming Human: A Theory of Ontogeny*. The Belknap Press of Harvard University Press.
- Tomasello, M. (2010). *Origins of Human Communication*. Bradford Books.
- Unnsteinsson, E. (2019). Silencing without Convention. *Pacific Philosophical Quarterly*, 100(2), 573–598. <https://doi.org/10.1111/papq.12260>
- Vesga, A. (forthcoming). Communicating Testimonial Commitment. *Ergo: An Open Access Journal of Philosophy*
- Williams, B. (2004). *Truth and Truthfulness: An Essay in Genealogy*. Princeton University Press.

CHAPTER 3

The Dual Nature of Testimony

An adequate theory of testimony must not only i) explain why testimony justifies belief and ii) explain why testimony entails second-personal relations of accountability. It shall also answer iii) why these two facts come together; it should, that is, account for a common explanatory link between them. Call it the Dual Nature Desideratum. This desideratum, I argue, reframes the challenges that traditional evidentialist and assurance views in the epistemology of testimony pose to each other. Correspondingly, here I defend a distinctive account of testimony that answers (i)-(iii). Cooperative action, I contend, offers a unified grounding source of the second-personal and evidential natures of testimony. Cooperative engagement between speaker and addressee enables our testimonial practices. In turn, the practical pressures of cooperative action provide distinctive epistemic grounds that justify belief and entitle audiences to hold speakers accountable for failing to fulfill their cooperative commitments.

1. Overview

Testimony is a ubiquitous epistemic source. Through the testimony of other people, I have learned that Simón Bolívar died in Santa Marta, the local farmers market sells plum tomatoes, and my blood type is A+. When things go well, the testimony of others justifies our beliefs.¹ Correspondingly, an adequate theory of testimony should explain why, or in virtue of what, testimony justifies belief. So, more schematically, we can pose the following desideratum:

Epistemic Desideratum: An adequate theory of testimony should offer an account of that in virtue of which testimony of p can justify believing p .

Let us say that whatever fulfills the ‘that’ slot in the desideratum will count as an explanation of the *epistemic nature* of testimony.

Yet, along with its epistemic nature, testimony has a distinctive normative aspect that other epistemic sources lack. If it turns out that my blood type is not A+, I can appropriately hold the source of my belief accountable for providing defective testimony. At the very least, I can appropriately question them on the reasons they had to offer defective testimony and reproach them if they offer no good excuse for doing so. Arguably, then, a complete theory of testimony should also answer why or in virtue of what testimony entitles these relations of accountability. Again, more schematically, we may pose a corresponding desideratum:

Accountability Desideratum: An adequate theory of testimony should offer an account of that in virtue of which testimony of p entitles its addressee to hold the testifying speaker accountable for the epistemic merit of her testimony.

¹ My aim here is to be as ecumenical as possible. Hence, ‘justification’ throughout the paper is neutral between externalist and internalist (and everything in between) conceptions of justification (cf. Bonjour, 1980; Bonjour & Sosa, 2003; Feldman & Conee, 2001; Goldman, 1988, 2009; Lasonen-Aarnio, 2010; Pritchard, 2012; Sosa, 2009; Williamson, 2000).

Let us say in turn that whatever fulfills the ‘that’ slot in the desideratum will count as an explanation of the *second-personal nature* of testimony.

Crucially, in as much as our concept of testimony validates intuitions about its epistemic and second-personal natures, it also validates the intuition that they are deeply related. Epistemically defective testimony entitles second-personal claims of accountability. Likewise, being accountable to whom we testify seems to be an essential part of our epistemically significant testimonial practices. An adequate theory of testimony, therefore, should provide an account of the relationship between the epistemic and the second-personal natures of testimony. Put otherwise, such a theory should account for the relationship between the grounds of testimony’s epistemic nature and the grounds of its second-personal nature. Hence, we may pose a third desideratum for a theory of testimony:

Dual Nature Desideratum: An adequate theory of testimony should offer an account of that in virtue of which the evidential and the second-personal natures of testimony *come together*.

Let us say that whatever fulfills the ‘that’ slot in the desideratum will count as an explanation of the *dual nature* of testimony. The task of fulfilling this desideratum will be my focus here.

The Dual Nature Desideratum, I contend, reframes the strongest challenges that two important families in the literature on testimony pose to each other. Classic evidentialist accounts of testimony, paradigmatically represented by reductionism and anti-reductionism, offer explanations of the justificatory grounds of testimony that stand oddly disconnected from its second-personal nature. In turn, assurance views fail to account for the epistemic nature of testimony by attempting to ground it on relations of assurance or responsibility. Hence, both these models fail to account for the Dual Nature Desideratum.

My central aim is to offer an account of testimony that fulfills this dual explanatory role. Following key insights in the philosophies of action and language, I argue that testimony is characterized by the practical pressures of cooperative action. To testify is to issue an utterance that must be interpreted as cooperatively offering information as means to achieve the shared goals that structure a given conversation.

The cooperative view synthesizes the lessons from evidentialist and assurance views of testimony. Testimony can only fulfill its justificatory function once the assumption that interlocutors are actually cooperative is discharged by relevant evidence. Yet, once the assumption of cooperation has been discharged, testimonial action yields a distinctive kind of justification with respect to testified contents. This kind of justification follows from the fact that cooperatively providing information instrumentally requires a distinctive kind of agential monitoring of the epistemic merit of one’s testimony. In turn, as in any other cooperative engagement, failure to fulfill the standards set by accepted shared goals amounts to an

instrumental failure to do what one has committed to do *with* others. This cooperative failure entitles second-personal accountability. Said accountability, however, does not constitute the grounds of the epistemic nature of testimony; it only shares these grounds.

In the next section, I say more to motivate the Dual Nature Desideratum, as it is not met by traditional evidentialist (§2.1) and assurance views (§2.2). In §3, I present my positive view of testimony as enabled by the rationality of cooperative action. In §4, I detail how this account of testimony explains why it justifies belief and how it differs from other similar accounts. In §5 I turn to how cooperation entitles second-personal accountability. I offer concluding remarks in §6.

2. Justification and Accountability

The relation between the epistemic and second-personal natures of testimony has been the (*de re*) locus of intense debate between what has been called assurance views and evidentialist views in the epistemology of testimony. This debate has extended for over two decades, and an exhaustive overlook of its turns is out of the scope of this paper. For my purposes, however, it will be useful to consider how the Dual Natures Desideratum reframes the challenges each of these views pose to the other. This will illuminate the success criteria for the view I begin developing in §3.

2.1. *Justification, not Accountability*

The Epistemic Desideratum is the focus of a wide slew of theories in the epistemology of testimony. Unsurprisingly, epistemologists of testimony are interested in accounting for the epistemic nature of testimony. The problem, some theorists argue, is that epistemically focused views fail to capture the essential second-personal nature of testimony (Hinchman, 2005; McMyler, 2011; Moran, 2018). If this point is to amount to an effective complaint, however, it must be parsed carefully.

Note, to begin with, that there is a normative distinction between how testimony justifies belief and how other acts justify belief.² Say Roger and Felipe are having a serious conversation. Roger tells Felipe, “I have a cold.” Felipe therefore acquires the belief that Roger has a cold. This is a paradigmatic case of testimony. If it turns out Roger does not have a cold, Felipe is entitled to ask him for the reasons he had for testifying and reproach him if these reasons do not amount to a good excuse for offering defective testimony.³

² There is another important, addressee-centered aspect of the second-personal nature of testimony. When a speaker testifies (at least in some cases), it is inappropriate for an audience to outright deny or refuse to believe in her testimony without providing good reasons for doing so. In Hinchman's (2005) words, if so, addressees “slight” their testifying counterparts. While not my focus, the view I develop here explains this as well. See n. 31.

³ Most obviously, defective testimony amounts to false testimony. However, testimony can be defective in that it fails to meet other warrant conditions. See n. 24.

This distinctive entitlement would not ensue if Felipe came to believe that Roger had a cold because, in the middle of their conversation, Roger took cold medicine or presented cold symptoms. In these cases, Roger's actions (or aspects of them) justify Felipe's belief that Roger has a cold. Yet, Felipe is not entitled to hold Roger accountable if his belief turned up false—certainly not in the way he would be if Roger had testified that he had a cold.⁴

Arguably, a complete theory of testimony would explain why, even though testimonial and non-testimonial acts are sources of evidence, only the former carry normative entitlements not entailed by any other evidential source. This, however, is no trivial task. The distinctive accountability of testimony does not amount to the boilerplate accountability entailed by mere intentional action. If Roger intentionally led Felipe to believe he had a cold by taking medicine or making his symptoms conspicuous, Roger would not entitle Felipe to hold him accountable in the way characteristic of testimony. Even if Roger were (intentionally) pretending to have a cold, he would still not be accountable in this way. In these cases, Roger's acts do not involve the commitment to a content that characterizes testimony and entitles Felipe's reproach. Evidence of this is that, although Roger would plausibly be accountable in some other ways, he can still validly deflect testimonial accountability ("I never told you I had a cold!").

Crucially, evidentialist views are orthogonal to this normative dimension. Consider the now classic debate between reductionists and anti-reductionists. Reductionists hold that the rationality of testimony reduces to the rationality of other epistemic sources (Fricker, 1994, 2004, 2006; Goldman, 1999; Shogenji, 2006). Felipe is justified in believing that Roger has a cold because he has good non-testimonial grounds to believe that Roger's testimony is linked to the truth. Yet, acting in a way that yields inferential grounds for belief does not entitle relations of accountability. Roger's taking cold medicine can give good inferential reasons for Felipe's belief. Even so, Felipe is not entitled to hold him testimonially accountable if his belief turns out false.⁵

A parallel issue lures classic anti-reductionism (e.g., Burge, 1993, 1997; Coady, 1994, 2000). This view holds that testimony is a *sui generis* epistemic source. Just as with perception and memory, Roger's testimony offers *a priori* justification for Felipe's belief. Nevertheless, there is nothing about *a priori* justification itself that explains testimonial accountability. Felipe's perception of Roger's symptoms may provide *a priori* justification to believe he has a cold. Again, however, Felipe is not entitled to hold Roger accountable if he is not sick.⁶

⁴ As I use it, the notion of entitlement refers to the normative entitlements to reproach an agent for their actions. Hence, it differs from epistemic entitlements as discussed by Burge (2018) and Graham (2016). Nevertheless, if the view I defend here is correct, epistemic and normative entitlements have a common source.

⁵ This issue extends to views that aim to other reductivist views (cf. Friedman, 1987; Goldman, 1999; Jones, 1999; Kusch, 2002; Keren, 2007; Lipton, 2007; Malmgren, 2006; Schiffer, 2006).

⁶ This issue extends to other anti-reductivist views (cf. Foley, 2001; Gibbard, 1992; Schmitt, 2002).

This is often presented as a core failure of reductionism and anti-reductionism (Moran 2018; McMyler 2012). The fact that classic accounts of the epistemic nature of testimony fail to characterize testimony's distinctive normative entitlements must mean that their accounts of how testimony justifies belief is internally defective.

Nevertheless, as several commentators have rightly noted, there is nothing inherently incoherent about providing a purely epistemic account of the justificatory grounds of testimony and expecting that its second-personal nature is explained by independent, non-epistemic grounds (cf. Goldberg, 2018; Schmitt, 2010). Perhaps that is a task for psychology or ethics (Lackey, 2008). In our terms, if all the epistemologist is looking for is an answer to the Epistemic Desideratum, then accusing her of not answering the Accountability Desideratum misses the mark.

There is, however, a cogent charge in the vicinity. Presumably, theorists of testimony are interested in building a cohesive theory that explains the distinctive features of our testimonial practices. And we have seen that, as an epistemic source, testimony involves practices of accountability. These second-personal entitlements seem to be a core feature of testimony, just as its epistemic role of justifying belief. Therefore, a complete theory of testimony is not exhausted by accounting for the Epistemic Desideratum; it must also answer the Accountability Desideratum.

Yet, for starters, it is not clear how other proposed realms, like psychology or ethics, could offer independent answers to the Accountability Desideratum. The second-personal nature of testimony is essentially normative. It refers to what testifying agents should do when testifying and why do their addressees have the normative standing to reproach them when they fail these standards. Psychology, however, is not in the business of producing normative theories. Further, the normativity at stake is not necessarily moral. Roger can be accountable in the way relevant here without being morally blameworthy. Perhaps Roger's mild irresponsibility or hypochondria makes him say he has a cold when he does not. Felipe can still reproach for offering defective testimony him without this being a moral failure on Roger's part.

Moreover, it seems natural to think that the explanation of why testimony entitles second-personal accountability is connected to the explanation of why it justifies belief. Presumably, the facts in virtue of which testimony justifies belief are not fully independent from the facts in virtue of which testimony entitles second-personal accountability. Hence, we should expect that a complete account of testimony will not only explain why testimony provides reasons to believe and why it involves normative entitlements, but also why the way in which it provides justifies belief involves a distinctively second-personal normativity that other epistemic sources lack. In our terms, a complete theory of testimony is subject to the

Dual Nature Desideratum; it must explain why testimony's epistemic and second-personal natures come together.

Once the theoretical aim is shifted to the Dual Nature Desideratum, evidentialist views are bound to yield incomplete theories of testimony. They simply bypass the Accountability Desideratum. Of course, building such a theory may not be what epistemically focused theorists aimed to do in the first place.⁷ Yet, it is what we aim to do here.

Admittedly, it remains open that our best concept of testimony will ultimately eschew the intuitively deep connection with second-personal entitlements. Some philosophers, for instance, have motivated the claim that testimonial justification and second-personal entitlements do not necessarily come together. This judgement is taken to be substantiated by cases where speakers do not seem to be accountable to those who gain testimonial knowledge from their testimony—paradigmatically, eavesdropping audiences (Lackey 2010; Owens 2006). However, given that the positive view I develop below accommodates these cases while also accounting for the Epistemic, Accountability, and Dual Nature Desiderata, it has an upper hand over a view that only accounts for the epistemic nature of testimony.

2.2. *Accountability, not Justification*

In a broad stroke, assurance views aim to account for the epistemic nature of testimony in terms of its second-personal nature. When a speaker testifies, the argument goes, she assures her audience (or assumes responsibility for the fact) that she has the appropriate epistemic privilege with respect to the content of her testimony. This irreducibly second-personal act of assurance is supposed to entail a distinct kind justification to believe in the contents of a speaker's testimony. Prime examples of this strategy are Hinchman (2005), McMyler, (2011), Moran (2018; 2013), and Ross (1986).⁸

Assurance views are in an excellent position to address the second-personal nature of testimony. By definition, when I assure you that I will do φ , and thus assume responsibility for doing φ , I entitle you to hold me accountable if I fail to do φ . Therefore, if testifying consists in assuring that or assuming responsibility for bearing the appropriate epistemic privilege with respect to the content of our words, then by definition it entitles the second-personal relations of accountability whose explanation we are after.

⁷ Characterizing the current landscape of the epistemology of testimony as neatly divided by reductionists and anti-reductionists would be shortsighted at best. Ever since the dawn of this distinction, there have been elaborations, and combinations of these views (Goldberg, 2007, 2012; Graham, 2000, 2016; Lackey, 2008; Schmitt, 2006). I postpone addressing how other evidentialist views fail to meet this challenge.

⁸ There are other views that, while related and inspired by aspects of assurance views, lie at the margins of the assurance family (cf. Faulkner, 2011, 2015; Pettit, 1995). In virtue of brevity, I postpone an argument of how they too fail to meet the dual nature desideratum.

It is not so clear, however, that assurance theories offer a satisfactory explanation of testimony's epistemic nature. An account of the justificatory grounds of an epistemic source must provide a justification-yielding link between said source and some epistemic good. It must answer precisely why relying on said source leaves individuals epistemically better off—with, say, more true beliefs, reasonable beliefs, or knowledge, depending on your epistemological allegiances. In other words, the assurance theorist must explain what about assuming responsibility one's words grounds the epistemic, justification-yielding link to said content. How is it that testimony justifies belief *in virtue of* the act of assurance?

Nevertheless, the sole act of assuring another person of the truth of *p* does not clearly provide itself an explanation of why it is epistemically rational to believe *p*. The sole fact that I can hold someone accountable if what they tell me is false does not make it rational for me to believe what they say. As Lackey (2008), Schmitt (2010), and Goldberg (2012; 2018) have noted, accepting the assurance of another person seems epistemically rational only if we have epistemic justification to think that such assurance is relevantly connected to the truth of *p*.

Of course, assurance theorists admit that it would be irrational to simply trust every person that assumes responsibility for the truth of what they are saying. There must be some background conditions that, if obtained in a testimonial interaction, enable the epistemic function of acts of assurance and give testimony its justifying function.⁹ At this point, however, assurance theories face a dilemma.¹⁰

On the one hand, the assurance theorist could characterize these background conditions in terms of another epistemic source that justifies the audience's belief in what is being assured. For instance, we could evaluate what independent reasons there are to think our testifying counterpart is honest and competent. Nevertheless, on this approach, the assurance itself is no longer grounding the epistemic grounds of testimony. As this path to explain the epistemic significance of testimony simply bypasses the very act of assurance.

On the other hand, assurance theorists could double down and argue that the background conditions that make it epistemically rational to believe a testifying speaker must be explained in terms of something irreducibly second-personal about the act of assuming responsibility for what they say. This path, however, faces the same issues that I have presented above. It is not clear what about relations of responsibility themselves connects testimony with an objective epistemic good. And without an irreducibly second-personal criterion that distinguishes between justifying and non-justifying assurances, the assurance project seems epistemically mute.¹¹

⁹Faulkner (2016) discusses this prospect. See also McMyler (2012).

¹⁰Here I follow Lackey (2008, ch. 10).

¹¹A similar argument is explored more fully by Lackey (2008) and Schmitt (2010).

The assurance theorist may be tempted to opt out. For instance, Richard Moran (2021) complains that his critics accuse him of not doing something he never aimed to do. He suggests that the question about how audiences can gain testimonial justification from what a speaker says is not interesting. Hence, his assurance view deals with other more interesting aspects concerning how agency distinguishes our testimonial practices. I agree that critics often misread Moran's project. Yet, Moran is not clear about why the epistemic question should not be asked of testimony while it can be asked of any other epistemic source.¹² Insofar we think that testimony justifies belief, we are warranted to ask what grounds the epistemic link between the act of assurance and the truth. Opting out of this question will simply yield an incomplete theory of testimony. It will not account for its epistemic nature and its relationship with the distinctive agential features that Moran sees in our testimonial practices.

3. Cooperative Interpretation

In what follows, I put forward a view that offers a common grounding source for the epistemic and the second-personal natures of testimony. My claim is that cooperative coordination enables testimony and thus grounds its distinctive epistemic and second-personal natures. In this section, I provide a characterization of testimony in terms of cooperation. In §4, I turn to show how it meets the Epistemic Desideratum, and in §5 to how it meets the Accountability Desideratum.

Let us begin with a Gricean insight: The non-language-specific nature of cooperation enables interlocutors to issue and interpret signals that inform each other about the world. This is the crux of Grice's (1968) Cooperative Principle, and it lies at the core of a plethora of theories of communication across disciplines.¹³ Hence, to see how cooperation characterizes testimonial communication, it will be useful to see how, in general, cooperative action allows us to interpret our counterpart's actions.

First, note that cooperation occurs when agents strive towards a shared interest or goal. Two (or more) agents have a *common goal* when they aim to bring about the same state of affairs. Two (or more) agents have a *shared goal* when they aim to bring about the same state of affairs together.¹⁴ My partner and I have the common goal of fixing the car if we both aim to have it fixed. We share this goal when we both accept to do so in a mutually responsive way

¹²Moran briefly suggests that this is because testimonial justification will vary according to the situation (2021, p. 726). Unless one thinks this variation is unexplainable, the search for an epistemic account explaining this variation is a genuine one. See §4 for such explanation.

¹³ See, for instance, Asher, 2013; Asher & Lascarides, 2013; Cohen, 1982; Cohen & Perrault, 1979; Grice, 1991; Lewis, 1975, 1980; Litman & Allen, 2003; Roberts, 2012, 2018; Simons, 2018; Stalnaker, 1984, 2014.

¹⁴ Here I am following Bratman's (1992, 1993, 2014) classic treatment of joint intentions. But see also Cohen & Levesque, 1991; Gilbert, 2009; Pacherie, 2012, 2013; Pollack, 1992. Note that this does not require, but is compatible with it, positing ontologically distinct collective beliefs as in Gilbert & Priest (2013). It is also compatible with a reductivist picture where all that is at stake is the iteration of mutually referring individual beliefs, *a la* Bratman (2014).

that requires us to coordinate and interlock our actions (Bratman 1995, 2014; Gilbert 2009; Pacherie 2012, 2013). If so, we are cooperating to fix the car.

Crucially, acting towards a shared goal enables a distinct kind of rationalization of an agent's actions. We can make sense of why a cooperating agent's action not only because she is committed to a goal but also because she is committed to achieving that goal *with* another agent. Say that it is public to my partner and I that the toolbox is in the garage. Say too that my partner goes to the garage. Given my belief that he is committed to fixing the car with me, I can explain his actions as directed to achieve this goal and thus conclude that he went for the toolbox and not to rest, for instance.

Two considerations substantiate this conclusion: i) given that my partner is committed to fixing the car, resting would contravene his goal; and, crucially, ii) given that he is committed to fixing the car *with me*, had he not gone for the toolbox, he would have made this clear. This second consideration follows from the fact that, given our common ground and our shared commitment to a goal, my partner can himself conclude that I will interpret his actions as his going for the toolbox. In turn, under the mutual understanding that I would explain his actions as going for the toolbox and, had he not gone for the toolbox, he would have told me so on pains of dissolving the shared nature of the goal to which he is committed.

Hence, my interpretation follows a distinctive interpretive pattern enabled by my partner's cooperative commitments. That he went for the toolbox is required by the best explanation of his actions as embedded in our cooperative commitment and given our common ground. When a proposition is required by the best explanations of our cooperative counterparts, given our common background information, call it a *propositional cooperative interpretation*.¹⁵ If I do generate a corresponding belief through such a rationalization, it will be a *doxastic cooperative interpretation*.

How does this pattern of action interpretation play into testimony? To cooperatively interpret our counterpart's actions as means to achieve our shared goals, we often must interpret them as meant to *actively inform* us about the world.¹⁶ If I utter the words "it's in the toolbox," that utterance is interpretable as providing relevant information because my doing so is required by the best explanation of my action as guided by our shared goal.¹⁷ That is, my partner can interpret my actions as a testimonial utterance: an utterance aimed to inform him

¹⁵ The quality of an interpretation is given by considerations of practical normative pressures, of coherence, stability, and simplicity of an intentional action, given and avowed goal, see Bratman, 2014; Bratman et al., 1988; Harman, 1990.

¹⁶ So construed, the view developed here is in line with John Greco's (2021) social account of the transmission of knowledge. See §4 for critical differences.

¹⁷ Some insinuations may fulfill a cooperative role and fail to be testimony. Yet, if an insinuation is to preserve its distinctive deniability (Camp 2019; Pinker et al 2008; Vesga ms), it must be merely *allowed* but not *required* by the best cooperative explanations. That is, it must be compatible with other contravening interpretations.

about the world.¹⁸ Crucially, cooperative interpretations must hinge on public information. If only I knew the toolbox in question, it would be a procedural mistake on my part to expect you to interpret my utterance as referring to the toolbox I have in mind. Against a public backdrop of information, interlocutors can produce rich cooperative actions whose aim is to be interpreted as informing about the world.¹⁹

Given this insight, we can say the following:

Testimony: By issuing an utterance U, a speaker S testifies a proposition p to her addressee A whenever, given information public to S and A, the best interpretation(s) of how U serves as a means towards the conversation's accepted shared goal require interpreting U as informing that p.

Testimony is the outcome of communicative interaction that hinges on cooperation. Under the assumption that they are engaged in cooperation, S produces a signal that A interprets as aimed to inform that p to achieve the goals to which S and A are jointly committed.

To clarify this position, it will be useful to consider the following worry. Speakers can issue deceiving testimony and testimonial interactions can occur in contentious, dubious, or coerced contexts. In these cases, it may seem that testimony ensues in the absence of cooperation and, hence, that cooperation is not essential to testimony.

Note, to begin with, that the notion of cooperation at work is permissive. All it takes for two or more agents to cooperate is that they present themselves as motivated towards a shared goal and act mutually and responsively to achieve it.²⁰ In this sense, whether agents cooperate is independent from the reasons that they have to do so. Perhaps my partner is helping me to fix the car because of our friendship. Perhaps, however, he is motivated to help me because of the effective threat of what I would do if he does not. In both cases he is cooperating with me to fix the car. Hence, *pace* Hornsby (1994) and Moran (2021), agents can coerce or manipulate others into cooperating towards their own goals. Likewise, testimony can be given in less-than-ideal cooperative engagements. If an utterance is interpreted as

¹⁸ Defining what actively informing amounts to is admittedly a difficult task. If the action is clear enough to call for the intention to communicate a sufficiently specific proposition as a cooperatively required interpretation, it can serve to testify. This characterization goes against the suggestions of philosophers like Fricker (2012) and Goldberg (2007; 2009), who argue that speakers can only testify to the literal content of their utterances. See Hawthorne (2012), Peet (2016), and Vesga (ms) for arguments against this suggestion.

¹⁹ Note that in focusing on cooperation, and not on speaker intention, avoids Moran's (2018) worries of rendering testimony too dependent on a one-way direction between speaker and addressee.

²⁰ Presenting oneself as interested in coordinating towards a shared goal may be as easy as simply issuing an utterance in the right context, continuing a conversation, or nodding while talking.

intended to inform under a shared goal, even if it is motivated by non-ideal reasons, said utterance will still count as testimony.²¹

Further, agents can accept a shared goal that differs from their personal goals. This makes sense of the fact that agents can accept shared goals in bad faith. Deceiving speakers can deviously testify by issuing utterances that are cooperatively interpreted as aimed towards the shared goals they have accepted but are not genuinely committed to.²² Nevertheless, if deception is to be successful, speakers must offer good (even if ultimately wrong) indication that they are indeed cooperatively engaged and are genuinely interested in providing information to their audience. By the same token, when there is doubt that a speaker is indeed committed to a shared goal she accepts, her utterance may still count as testimony if it plays a cooperative role with respect to said goal. Yet, we will see, under these conditions, testimony will not fulfill its epistemic function. I say more about how contentious and dubious contexts affect the epistemic and second-personal aspects of testimony in the next sections.

Before that, one final clarification. It is important to distinguish between addressees and audiences. An addressee is that to whom the speech act of testimony is directed and with whom the speaker must share a goal to testify.²³ An audience is whoever witnesses testimony. I have said that testimony is enabled by the cooperative interaction between a speaker and an addressee. However, below I argue that audiences are justified to believe in the cooperatively justified contents of a speaker's testimony. This distinction will not matter in several cases, as audiences and addressees often coincide. Nevertheless, it will be crucial to explain cases in which agents can gain testimonial from others with whom they share no goals.

4. Cooperatively Justified

To account for the epistemic nature of testimony, the cooperative view should explain how the fact that it is practically rational to provide information as cooperative means towards a shared goal can make it epistemically rational to believe said information. Cooperative reasoning not only rationalizes my act of testifying that the wrench is in the toolbox; it also justifies your belief that the wrench is in the toolbox.

²¹ This already differentiates the present view from theories that are too optimistic about testimony. McMyler (2011), for instance, states that testimony essentially involves some kind of “goodwill” from the speaker to the audience. Hinchman (2014) argues that the epistemic import of testimony hinges on a “sympathetic” concern testifying speakers have towards the epistemic concerns of their audiences. Similarly, Greco (2021) argues that testimony has a distinctive epistemic function only when the interlocutors’ motivations fully align.

²² The difference between personal and shared goals also explains how speakers can testify accidentally. Once a speaker has accepted a shared goal and issued an utterance that requires a cooperative interpretation, she has testified to it—even if this was not her privately avowed intention.

²³ Addressees may be unknown to the speaker, as when I talk to someone behind a curtain. They may also be potential or located in the future, as when I leave a note saying that the office fridge is out of order. Moreover, addressees may be groups of people, as when I say to the guests that dinner is ready.

To begin, note that successful testimony ensues when the epistemically significant action of providing an epistemic good serves as means towards an accepted shared goal. Your having the information that the wrench is in the toolbox is a means towards fixing the car. Crucially, failing to have the epistemic good one offers in testifying amounts to an instrumental failure to achieve said shared goal. Put otherwise, from the perspective of the speaker, testimony involves an epistemically significant instrumental standard, stemming from cooperative action: if I am to cooperatively offer an epistemic good (i.e., to testify), I must have said good. Here I bracket the issue of what specific epistemic good this norm requires from testifying speakers.²⁴ What matters for my purposes is that testimony involves epistemic standard that, on pains of instrumental failure, must be met to comply with its cooperative purpose.²⁵

Of course, the fact that testimony is subject to an epistemic norm does not mean that every instance of testimony complies with this norm. A speaker may present herself as cooperatively issuing an epistemic good but fail to be actually cooperative; she may, for instance, be dishonest or incompetent. Hence, from the audience's perspective, if testimony of a speaker is to fulfill the epistemic role of providing relevant information, the if-clause of the instrumental norm must be discharged. It must be the case that the speaker is actually—and not just presenting herself as—in a position to cooperate (i.e., properly motivated and able to actually advance the purposes of the conversation.)

This means that testimonial justification is two-tiered. First, the assumption of cooperation must be discharged by evidence that the speaker is actually cooperative and is in position to comply with the instrumental norm. I remain neutral here on the kind of evidence that appropriately discharges the assumption of cooperation. We may be *a priori* entitled to believe that agents that present themselves as cooperative are actually cooperative or we may

²⁴ What I have said is compatible with testimony requiring, for instance, knowledge (Williamson, 1996, 2000), certainty (Stanley, 2008), justified belief (McKinnon, 2015), or a context-sensitive standard (Goldberg 2015, 2018). Let me, however, advance a tentative argument for a truth and an anti-luck components to the instrumental norm. Cooperative failure is most obvious when testimony is false. If your having the information that the wrench is in the toolbox is required to achieve our aims, then telling you that the wrench is in the toolbox when it is not will directly contravene our shared goal. Further, merely guessing is arguably a cooperative failure as well. Say you ask me to invite our mutual friend Ana to our planned dinner. Yet, I decide to achieve this goal by dialing a random ten-digit number. Given the extremely low chance that I will succeed, I failed to act cooperatively—even if by sheer chance I dial the right number. My act falls short of intentionally bringing about my cooperative commitment. By the same token, merely guessing when we testify amounts to a failure to actually be a cooperative interlocutor.

²⁵ This brings the cooperative picture of testimony close to norm-centered views of the nature of assertion (e.g., Williamson 2000, Stanley 2005, Goldberg 2018). I take it as a strength of the cooperative view that it can capture the clear theoretical benefits of these accounts of assertion (see Goldberg 2018 for a compelling argument in this respect). A key difference with other accounts is that the norm at stake is not game-like (Williamson, 2000), social (Graham 2020), or conventional (Kölbel 2016), but instrumental.

need independent evidence that substantiate said belief.²⁶ Hence, the epistemic significance of discharging the cooperative assumption is congenial with evidentialist criteria, as long as they refer to our justification for believing that an agent is acting cooperatively.²⁷

Crucially, note—as this will be important shortly—that, most plausibly, the motivational structure in which a conversation is embedded has an important role in determining the evidence necessary to discharge the assumption of cooperation, no matter how this evidence is defined. If testimony ensues in contexts where cooperating agents are clearly motivated to achieve the same outcome, like two friends fixing a car to use it later, discharging the assumption of cooperation will not take much. If testimony ensues in contexts where there is *prima facie* indication that motivations do not fully align, like conversations with hostile interrogators or contentious negotiations, it will take more to ensure that offered testimony is indeed playing a cooperative role.

Nevertheless, the epistemic nature of testimony does not reduce to the epistemic significance of whatever evidence is required to discharge the assumption of cooperation. Once the assumption of cooperation has been discharged, testimony yields a distinctive pattern of justification. A speaker testifies when she presents her utterances as cooperatively offering information as means to achieve the goals that structure a conversation. If she is actually committed, offering defective testimony amounts to an instrumental failure to do what she committed to do. Hence, a genuinely cooperative speaker would not have informed that *p* unless she had adequate authority over *p*. Under the (discharged) assumption that the speaker is genuinely cooperative, the link between a piece of testimony and the truth is mediated by the very agency of the speaker and her commitment to cooperate with another agent. There is a connection between an actually cooperative speaker's testimony and the truth *because* she is committed to ensure so.

This second tier of testimonial justification distinguishes testimony from other epistemic sources. And, crucially, it is not present in cases where the practical reasons of agents do not involve the cooperative commitment to provide an epistemic good. Although Felipe can learn that Roger has a cold by inferring this from the fact that he took cold medicine, this belief is not justified by Roger's commitment to providing information about which he has an epistemic privilege. In contrast, when Roger testifies that he has a cold, and he is indeed

²⁶I am inclined to be a pluralist. We are *prima facie* entitled to believe that our seemingly cooperative interlocutors are indeed genuinely cooperative, given two interacting considerations: a) our cognitive abilities are enabled, fostered, and evolved by our species' ability to default into actions that yield cooperation and large-scale coordination (Henrich, 2016; Tomasello, 2015); and b) societies generate social norms as effective means to externally motivate interlocutors cooperate (cf. Bicchieri, 2006; Graham, 2020a, 2020b; Starr, 2018). Yet, the assumption can be independently discharged when these considerations are defeated.

²⁷Therefore, this kind of evidence pertains to a wider realm than testimony. It is the same kind of evidence that warrants our believe that our cooperating partners are doing their part in the joint project of fixing the car.

committed to being a cooperative speaker, he would not have easily testified that he has a cold unless he had the corresponding epistemic authority over his testimony. As opposed to other epistemic sources, cooperative agency is testimony's distinctive epistemic difference maker.

We can say, then, that the evidence required to discharge the assumption of cooperation *enables* testimony's distinctive epistemic significance. When things go right, testimonial belief is justified by the speaker's agential monitoring of the epistemic merit of what she says, given that this monitoring is required by the practical structure of what she has cooperatively committed to do with her audience. Thus, testimony has its epistemic character in virtue of the justificatory pattern enabled by cooperative reasoning. In a slogan, actual (discharged) cooperation provides a secure translation key between the speaker's reasons to act and the audience's reasons to believe.

At this point, it will be helpful to note a critical difference with John Greco's (2021) neighboring account. Greco and I hold that cooperation gives testimony its distinctive epistemic profile. Yet, For Greco, there is a fundamental epistemic difference between the following two cases: i) your friend tells you he is going to be at your party and ii) a job applicant tells her potential employer that she has relevant previous employment.²⁸

In (i), your friend is motivated to cooperate and thus to relieve you from your epistemic duties. Therefore, Greco argues, she does all the epistemic work necessary for you to gain the corresponding justification and you can fully epistemically depend on her. In this case, testimony serves an irreducible epistemic function. In (ii), however, the applicant may not be so motivated. Hence, the employer has to appeal to further information to make sure the applicant is not lying. So, the argument goes, the epistemic significance of the applicant's testimony reduces to that of the evidence the employer appeals to in order to make sure the applicant is not lying. Therefore, only in the first case testimony counts as an epistemically distinctive source of justification.

In contrast, on the view presented here, the perceived difference between these cases owes to the fact that, given the motivational structure of an interaction, it will take more or less to discharge the assumption of cooperation. In both cases, the testifying speakers present themselves as cooperative interlocutors and the justifying potential of their interaction depends on whether they indeed are cooperative interlocutors. Hence, in both cases, the assumption of cooperation must be discharged. Yet, the motivational structure of your interaction with your friend easily discharges the assumption. In the second case, however, it will take more to discharge the assumption that the applicant is not presenting herself as cooperative in bad faith. Perhaps, for instance, the employer must make sure the applicant is

²⁸ These are Greco's examples, 2021: p.33.

properly motivated to cooperate by the detrimental consequences of providing false information.

Nonetheless, once the assumption of cooperation is discharged, testimony serves the same epistemic role in both cases. Both your and the employer's beliefs will be justified by the fact that both speakers are indeed committed to cooperate and, on pains of instrumental failure, would not have done so unless they had good reason to believe what they say. In both cases, even if the motivations to cooperate are different, the epistemic significance of their testimony hinges on their cooperative commitment to monitor the epistemic merit of their words.

Crucially, Greco's view misses the fact that there is a range of cases between robust and contentious cooperation. If there is indeed a difference of epistemic kind between these two extremes, there must be a principled threshold that distinguishes which middle cases testimony counts as a distinctive epistemic source and in which it does not. Yet, Greco has not provided such criterion. In contrast, under the cooperative view presented here, the degree of cooperation will modulate how much it will take to discharge the assumption of cooperation and thus enable a given instance of testimony to fulfill its distinctive epistemic role. Hence, the difference between the cases at stake is one of degree, not of kind.

Another key aspect of the view I present here is that it accommodates eavesdropping cases that put pressure on assurance views and Greco's notion that one-directional trust is the epistemic difference maker in testimony (Lackey 2010; Owens 2006; McMyler 2011). Say Fiona testifies to Paulo that there is lager on tap at the Easty Bar while, unbeknownst to them, Mischa is eavesdropping. Arguably, if Paulo can gain testimonial justification from this, Mischa can too. Yet, Fiona is not making herself responsible to Mischa nor is Mischa's epistemic dependence playing any role in Fiona's reasoning. So, it seems that direct second-personal trust or responsibility are not essential to explanation of why Mischa is testimonially justified to believe that there is lager on tap at the Easty Bar.

Nevertheless, as long as Mischa understands the epistemic significance of the cooperative interaction between Fiona and Paulo—as long as she understands that Fiona is genuinely committed to acting cooperatively towards a goal shared with Paulo and how this goal entails epistemic standards—she has access to the cooperative justification that characterizes testimony.²⁹ She knows Fiona is committed to ensuring the relevant epistemic link between her testimony and the truth. While there must be a cooperative relationship

²⁹ This will also accommodate why we may gain testimonial knowledge from diaries or soliloquies (Lackey 2010; Owens 2006). In as much as in these cases the speaker can be understood as cooperating with her future time slices, an utterance in a diary may count as testimonial. If I read Lee's diary and know he uses it as a log for his house chores, I can gain testimonial knowledge that there is a spider nest in his attic. Otherwise, decontextualized diary entries and soliloquies do not count as testimony. This should not be too high of a cost to bear. For we may still gain knowledge from them, just not testimonial knowledge (cf. Coady 2000; Moran 2018).

between speaker and addressee for testimony to ensue, testimony can provide justification to audiences who are not embedded in this relationship.

To sum up, testimonial justification is enabled by evidence that the speaker is properly motivated and able to pursue the shared goal that structures the conversation. Hence, this account of testimonial justification is amicable to evidentialist views. Nonetheless, as the assurance intuition has it, testimonial justification is still distinctively second-personal. Once the assumption of cooperation is discharged, the epistemic function of testimony is the outcome of a rational interaction that can only ensue through the mutual recognition of conversing interlocutors that they are committed to a shared goal. We have, then, a distinctively second-personal and genuinely epistemic answer to the Epistemic Desideratum. I turn now to how this same practical structure yields an answer to the Accountability Desideratum.

5. Cooperatively Accountable

Testimonial justification ensues only if offering an epistemic good serves as means towards a goal jointly accepted by a speaker and her addressee. From this, we have seen, it follows that defective testimony amounts to a failure to act towards a goal to which one has committed *with* another agent. This kind of failure to act towards an avowed shared goal entitles agents to react in ways that monitor and correct cooperative behavior—crucially, holding each other accountable for failing to comply with the requirements of a shared goal.

If I fail to achieve a goal you and I are jointly committed to, I am not only failing to act according to my own commitment, but I am also failing *you*. In the right circumstances—i.e., barring genuine excuses—this failure entitles you to appropriately hold me accountable.³⁰ If you and I are committed to painting the whole house white, you can genuinely criticize me if I paint the north wall green.

Correspondingly, testifying speakers who lack the epistemic privilege required by the instrumental standard of testimony betray the commitment that enabled them to testify in the first place. This betrayal entitles audiences to hold speakers accountable for the contents of their testimony.³¹ My telling you that the wrench is in the toolbox as a cooperative means to our goal entitles you to hold me accountable to at least explain myself if the wrench is not in the toolbox. Therefore, once the utterances of a testifying speaker yield a required cooperative interpretation, they also entitle the addressee to hold the speaker accountable when the utterance's epistemic import is defective.

³⁰ What matters is that once my counterpart has testified, she is accountable to offer a reasonable excuse for her failure.

³¹ As suggested in n. 2, this yields an explanation of that other addressee-centered normative aspect of testimony. Given that the audience is also committed to a goal that requires her accepting cooperatively offered information, at least sometimes, once a speaker has testified, it is a procedural failure on the audience's part to reject cooperatively offered information without good reason—and thus “slights” her counterpart (Hinchman's 2005).

Crucially, this explains why testimonial accountability is absent in the case that Felipe comes to believe Roger has a cold because he took cold medicine or presented cold symptoms. In these cases, Felipe's coming to believe Roger has a cold does not hinge on Roger's commitment to a shared goal that instrumentally requires him to provide good reasons to believe. Hence, even if he is pretending to have cold symptoms and it was his intention to generate a false belief in Felipe, his action does not amount to a betrayal of a cooperative commitment—as it would be had he actively testified he had a cold.

This conclusion is compatible with the intuition that speakers are not accountable to eavesdropping audiences in cases like the one above. Granted, Fiona may not be testimonially accountable to Mischa, who overheard her testimony to Paulo. However, Fiona was able to testify because she engaged in cooperation with Paulo. Fiona is therefore accountable to Paulo. Testimony entails relations of accountability, even if not with every single agent that can obtain justified beliefs from it.³²

This is one of the main perks of providing a common source for testimony's epistemic and second-personal natures. Cooperation enables the justificatory grounds of testimony and entails relations of accountability. Yet, *contra* the assurance theorist, these relations of accountability are not themselves providing the justificatory grounds of testimony. Eavesdroppers can still be justified when witnessing a testimonial engagement even if the speaker is not directly accountable to them. Shifts in the relations of accountability do not necessarily entail shifts in the distinctive epistemic merits of an instance of testimony.

One may still wonder about the status of this second-personal entitlement when speakers testify in bad faith. Deceiving speakers present themselves as cooperatively committed to a shared goal to which they are not, in fact, personally committed. Omar may utter "At 10:00am" to deceive Gale into falsely believing that her doctor's appointment is at ten in the morning. If his deception is to be successful, he must present himself as committed to a goal such that his utterance is cooperatively rationalized as providing said information. In that case, it may well be that Gale's claim of accountability has no internal purchase on Omar, given his private motivations.

Nonetheless, *within* the instrumental requirements set by the shared goals that enable testimonial engagement, Omar's deception is still an instrumental failure. He fails to comply with the cooperative commitment that enabled him to testify (albeit in bad faith). Therefore,

³² Is Lee from n. 29 accountable to himself? In a way, he is. When he reads his diary at t_2 , Lee can reproach himself at t_1 for giving him defective testimony. Admittedly, in terms of possible compensation for this failure, the only object of reproach is Lee at t_2 .

Gale is still entitled to her accountability claim; for it is exactly the kind of behavior that would preserve the goal that structures their conversation.³³

Importantly, the cooperative view defended here is not committed to equating testimonial accountability and moral accountability. This would require a substantive view on the relation between joint action and moral responsibility. I do not need to solve this issue here to note that, given its cooperative nature, testimony involves at least a sense of instrumental accountability.

Yet, independently of its moral import, the cooperative view offers a fulcrum to explain and model how the distinctive accountability of testimony is deeply connected to broader social phenomena. If I am correct, testimony ensues when interlocutors engage cooperatively towards a shared goal. Nonetheless, we have seen, cooperation can be motivated in several, and often non-ideal, ways. Powerful agents may coerce others into cooperatively acting towards their own individual goals. I may be motivated to help my boss move his couch because he has threatened to fire me if I do not. Further, communities often have external social norms and sanctions in place that motivate cooperation among agents who would otherwise fail to cooperate (Bicchieri, 2016).

If the present view is correct and the dual nature of testimony is grounded on cooperative action, then social phenomena that guide and modulate how and why we cooperate inherently shift the relations of accountability between interlocutors in a testimonial interaction. Cooperation serves as an explanatory fulcrum to account for how these phenomena—or, as Goldberg (2020) puts it, these conversational pressures—determine to whom and for what are we accountable in testimony. The view I advance here offers the prospects of a unified explanatory framework that captures not only the grounds of the second-personal and epistemic natures of testimony, but how these grounds are inherently sensitive to large-scale social phenomena, like social norms and social power (e.g., Brown & Levinson 2000; Brown 2010; Hesni 2012; Langton 2019; Schiller 2019; Goldberg 2020).

Although more needs to be said, this means that the kind of second-personal accountability of testimony can often go wrong in many cases. I may testify to my boss because I am fearful of the consequences of not doing so. In turn, my boss may be justified to believe my testimony if his threat is effective enough to motivate my cooperation. If so, he has successfully enacted his power in a way that enables testimonial justification based in my coerced commitment to cooperate. In turn, offering false information is still a (procedural)

³³ The notion that there is something wrong with committing to joint goals and then defecting or betraying them has wide intuitive appeal (Bratman, 2009, 2014; P. R. Cohen & Levesque, 1991; Gilbert, 2009; Lewis, 1975; Michael & Pacherie, 2015; Pacherie, 2013; Pollack, 1992). However, it must be the case that sometimes it is rational to defect from our cooperative commitments. My claim is that in as much as it is a failure to defect from a joint commitment, it is a failure to offer epistemically defective testimony.

failure on my part to comply with this commitment. Therefore, the second-personal normativity of testimony can be coopted by wider collective-level social factors—power relations in this case—that motivate agents into cooperation.

To sum up, *contra* the assurance theorist, the fact that agents are accountable for the epistemic merit of their testimony is not the reason why testimony has its epistemic significance. Instead, cooperative action constitutes a unified ground of these two natures. The fact that testifying speakers are accountable to their audiences is explained by the same grounds that explain the fact that audiences can be justified in believing a speaker's testimony. Thus, accountability is essentially related to what makes testimonial justification unique. Yet, accountability itself does not explain testimonial justification.

6. Conclusion

I have presented a desideratum for an adequate theory of testimony that two important families of views in the epistemology of testimony fail to meet. Extant views in the literature fail to provide a common explanation of why testimony's epistemic and second-personal natures come together. I have also presented a characterization of testimony as enabled by cooperative action that can meet the challenge. To testify is to cooperatively offer information. In turn, cooperation provides a common source for testimony's justificatory status and its second-personally normative entitlements. The epistemic and second-personal grounds of testimony stem from our ability to translate the practical requirements set by shared goals to epistemic considerations for our beliefs about the world. By providing a common explanatory link, this cooperative account meets the Dual Nature Desideratum.

References

- Asher, N. (2013). Implicatures and discourse structure. *Lingua*, 132, 13–28.
- Asher, N., & Lascarides, A. (2013). Strategic conversation. *Semantics and Pragmatics*, 6(0), 2-1–62.
- Bicchieri, C. (2006). *The Grammar of Society: The Nature and Dynamics of Social Norms*. Cambridge University Press.
- Bonjour, L. (1980). Externalist Theories of Empirical Knowledge. *Midwest Studies in Philosophy*, 5(1), 53–73. h
- BonJour, L., & Sosa, E. (2003). *Epistemic Justification: Internalism Vs. Externalism, Foundations Vs. Virtues*. Wiley-Blackwell.
- Bratman, M. E. (1992). Shared Cooperative Activity. *The Philosophical Review*, 101(2), 327–341. JSTOR.
- Bratman, M. E. (1993). Shared Intention. *Ethics*, 104(1), 97–113. JSTOR.
- Bratman, M. E. (2009). Modest Sociality and the Distinctiveness of Intention. *Philosophical Studies*, 144(1), 149–165.

- Bratman, M. E. (2014). *Shared Agency: A Planning Theory of Acting Together*. In *Shared Agency*. Oxford University Press.
- Bratman, M. E., Israel, D. J., & Pollack, M. E. (1988). Plans and resource-bounded practical reasoning. *Computational Intelligence*, 4(3), 349–355.
- Burge, T. (1993). Content Preservation. *Philosophical Review*, 102(4), 457–488.
- Burge, T. (1997). Interlocution, Perception, and Memory. *Philosophical Studies*, 86(1), 21–47.
- Camp, E. (2019). *Insinuation, Common Ground, and the Conversational Record*. In *New Work on Speech Acts*. Oxford University Press.
- Coady, C. A. J. (1994). Testimony, Observation and “Autonomous Knowledge.” In A. Chakrabarti & B. K. Matilal (Eds.), *Knowing From Words* (pp. 225–250). Kluwer Academic Publishers.
- Coady, C. a. J. (2000). *Testimony: A Philosophical Study*. In *Testimony*. Oxford University Press.
- Cohen, L. J. (1982). What Is Necessary for Testimonial Corroboration? *The British Journal for the Philosophy of Science*, 33(2), 161–164. JSTOR.
- Cohen, P. R., & Levesque, H. J. (1991). Teamwork. *Nous*, 25(4), 487–512. <https://doi.org/10.2307/2216075>
- Cohen, P. R., & Perrault, C. R. (1979). Elements of a plan-based theory of speech acts. *Cognitive Science*, 3(3), 177–212. [https://doi.org/10.1016/S0364-0213\(79\)80006-3](https://doi.org/10.1016/S0364-0213(79)80006-3)
- Cowie, C., & Greenberg, A. (2018). Constitutivism about Epistemic Normativity. In C. Kyriacou & R. McKenna (Eds.), *Metaepistemology: Realism & Antirealism* (pp. 173–196). Palgrave Macmillan.
- DeRose, K. (2011). *The Case for Contextualism: Knowledge, Skepticism, and Context, Vol. 1* (Reprint edition). Oxford University Press, U.S.A.
- Faulkner, P. (2007). On Telling and Trusting. *Mind*, 116(464), 875–902.
- Faulkner, P. (2011). *Knowledge on Trust*. Oxford University Press.
- Faulkner, P. (2015). The attitude of trust is basic. *Analysis*, 75(3), 424–429.
- Feldman, R., & Conee, E. (2001). Internalism Defended. *American Philosophical Quarterly*, 38(1), 1–18.
- Foley, R. (2001). *Intellectual Trust in Oneself and Others*. Cambridge University Press.
- Fricke, E. (1994). Against Gullibility. In B. K. Matilal & A. Chakrabarti (Eds.), *Knowing from Words: Western and Indian Philosophical Analysis of Understanding and Testimony* (pp. 125–161). Springer Netherlands. https://doi.org/10.1007/978-94-017-2018-2_8
- Fricke, E. (2004). Testimony: Knowing Through Being Told. In M. Sintonen, J. Wolenski, & I. Niiniluoto (Eds.), *Handbook of Epistemology* (pp. 109–130). Kluwer Academic Publishers.

- Fricke, E. (2006). Second-Hand Knowledge. *Philosophy and Phenomenological Research*, 73(3), 592–618.
- Friedman, R. (1987). Route Analysis of Credibility and Hearsay. *Yale Law Journal*, 96(4).
<https://digitalcommons.law.yale.edu/ylj/vol96/iss4/1>
- Gibbard, A. (1992). *Wise Choices, Apt Feelings: A Theory of Normative Judgment* (Reprint edition). Harvard University Press.
- Gilbert, M. (2009). Shared intention and personal intentions. *Philosophical Studies*, 144(1), 167–187.
- Goldberg, S. (2018). *Reasons and Testimony* (D. Star, Ed.; Vol. 1). Oxford University Press.
- Goldberg, S. C. (2015). *Assertion: On the Philosophical Significance of Assertoric Speech*. Oxford University Press.
- Goldberg, S. C. (2018). *To the Best of Our Knowledge: Social Expectations and Epistemic Normativity*. Oxford University Press.
- Goldman, A. I. (1988). Strong and Weak Justification. *Philosophical Perspectives*, 2, 51–69.
- Goldman, A. I. (1999). *Knowledge in a Social World*. Oxford University Press.
- Goldman, A. I. (2009). Internalism, Externalism, and the Architecture of Justification. *Journal of Philosophy*, 106(6), 309–338. <https://doi.org/10.5840/jphil2009106611>
- Graham, P. J. (2020a). The Function of Assertion and Social Norms. In S. Goldberg (Ed.), *The Oxford Handbook of Assertion* (pp. 725–748). Oxford University Press.
- Graham, P. J. (2020b). Assertions, Handicaps, and Social Norms. *Episteme*, 17(3), 349–363.
- Greco, J. (2008). What’s Wrong with Contextualism? *The Philosophical Quarterly* (1950-), 58(232), 416–436. JSTOR.
- Greco, J. (2010). *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity*. Cambridge University Press.
- Grice, P. (1991). *Studies in the Way of Words*. Harvard University Press.
- Harman, G. (1990). Reasoning, Meaning, and Mind. In *Reasoning, Meaning, and Mind*. Oxford University Press.
- Henrich, J. (2016). *The Secret of Our Success: How culture is driving human evolution, domesticating our species and making us smarter*. Princeton University Press.
- Hinchman, E. S. (2005). Telling as Inviting to Trust. *Philosophy and Phenomenological Research*, 70(3), 562–587.
<https://doi.org/10.1111/j.1933-1592.2005.tb00415.x>
- Jones, K. (1999). Second-Hand Moral Knowledge. *Journal of Philosophy*, 96(2), 55–78.
- Keren, A. (2007). Epistemic Authority, Testimony and the Transmission of Knowledge. *Episteme*, 4(3),
- Kusch, M. (2002). *Knowledge by Agreement: The Programme of Communitarian Epistemology*. Oxford University Press.

- Lackey, J. (2007). Norms of Assertion. *Noûs*, 41(4), 594–626.
- Lackey, J. (2008). *Learning From Words: Testimony as a Source of Knowledge*. Oxford: Oxford University Press.
- Lasonen-Aarnio, M. (2010). Unreasonable Knowledge. *Philosophical Perspectives*, 24(1), 1–21.
- Lewis, D. K. (1975). Languages and Language. In K. Gunderson (Ed.), *Minnesota Studies in the Philosophy of Science* (pp. 3–35). University of Minnesota Press.
- Lewis, L. (1980). *Convention*. Blackwell.
- Lipton, P. (2007). Alien Abduction: Inference to the Best Explanation and the Management of Testimony. *Episteme*, 4(3), 238–251.
- Litman, D. J., & Allen, J. F. (2003). Discourse Processing and Commonsense Plans. In P. R. Cohen, J. Morgan, & M. E. Pollack (Eds.), *Intentions in Communication*. The MIT Press.
- Malmgren, A.-S. (2006). Is There a Priori Knowledge by Testimony? *Philosophical Review*, 115(2), 199–241.
- McHugh, C., & Whiting, D. (2014). The Normativity of Belief. *Analysis*, 74(4), 698–713.
- McKinnon, R. (2015). *The Norms of Assertion: Truth, Lies, and Warrant* (2015 edition). Palgrave Macmillan.
- McMyler, B. (2011). *Testimony, Trust, and Authority*. Oxford University Press.
- Michael, J., & Pacherie, E. (2015). On Commitments and Other Uncertainty Reduction Tools in Joint Action. *Journal of Social Ontology*, 1(1), 89–120.
- Moran, R. (2018). *The Exchange of Words: Speech, Testimony, and Intersubjectivity*. Oxford University Press.
- Moran, R. A. (2013). I—Richard Moran: Testimony, Illocution and the Second Person. *Aristotelian Society Supplementary Volume*.
- Murray, S. E., & Starr, W. B. (2021). The Structure of Communicative Acts. *Linguistics and Philosophy*, 1–50.
- Owens, D. (2006). Testimony and Assertion. *Philosophical Studies*, 130(1), 105–129.
- Pacherie, E. (2012). *The Phenomenology of Joint Action: Self-Agency vs. Joint-Agency*.
- Pacherie, E. (2013). Intentional joint agency: Shared intention lite. *Synthese*, 190(10), 1817–1839. JSTOR.
- Pettit, P. (1995). The Cunning of Trust. *Philosophy and Public Affairs*, 24(3), 202–225.
- Pinker, S., Nowak, M. A., & Lee, J. J. (2008). The logic of indirect speech. *Proceedings of the National Academy of Sciences*, 105(3), 833–838.
- Pollack, M. E. (1992). The Uses Of Plans. *Artificial Intelligence*, 57, 43–68.
- Pritchard, D. (2012). Anti-Luck Virtue Epistemology. *Journal of Philosophy*, 109(3), 247–279.

- Roberts, C. (2012). Information structure in discourse: Towards an integrated formal theory of pragmatics. *Semantics and Pragmatics*, 5. <https://doi.org/10.3765/sp.5.6>
- Roberts, C. (2018). *Speech Acts in Discourse Context*. In *New Work on Speech Acts*. Oxford University Press.
- Ross, A. (1986). Why Do We Believe What We Are Told? *Ratio*, 1, 69–88.
- Schiffer, S. (2006). The Things We Mean. *Philosophy and Phenomenological Research*, 73(1), 208–210.
- Schmitt, F. F. (2002). Testimonial justification: The parity argument. *Studies in History and Philosophy of Science Part A*, 33(2), 385–406. [https://doi.org/10.1016/S0039-3681\(02\)00011-0](https://doi.org/10.1016/S0039-3681(02)00011-0)
- Schmitt, F. F. (2010). The Assurance View of Testimony. In A. Haddock, A. Millar, & D. Pritchard (Eds.), *Social Epistemology* (pp. 216–242). Oxford University Press.
- Shah, N., & Velleman, J. D. (2005). Doxastic Deliberation. *Philosophical Review*, 114(4), 497–534. <https://doi.org/10.1215/00318108-114-4-497>
- Shogenji, T. (2006). A Defense of Reductionism about Testimonial Justification of Beliefs. *Noûs*, 40(2), 331–346.
- Simons, M. (2018). Convention, Intention, and the Conversational Record. In *Beyond Semantics and Pragmatics*. Oxford University Press.
- Sosa, E. (2009). A Virtue Epistemology: Apt Belief and Reflective Knowledge, Volume I. *Analysis*, 69(2), 382–385.
- Stalnaker, R. (1984). *Inquiry*. Cambridge University Press.
- Stalnaker, R. (2014). *Context*. Oxford University Press.
- Stanley, J. (2005). *Knowledge and Practical Interests*. Oxford University Press.
- Stanley, J. (2008). Knowledge and Certainty. *Philosophical Issues*, 18(1), 35–57.
- Starr, W. B. (2018). *Socializing Pragmatics* (Vol. 1). Oxford University Press.
- Tomasello, M. (2015). *A Natural History of Human Morality*. Harvard University Press.
- Velleman, D. (2000). *The Possibility of Practical Reason* (Vol. 106, Issue 4, pp. 694-). Oxford University Press.
- Vermaire, M. (2020). Retraction and testimonial justification: A new problem for the assurance view. *Philosophical Studies*, 177(12), 3959–3972. <https://doi.org/10.1007/s11098-020-01417-y>
- Vesga, A. (manuscript). *Talking Together: On the Normativity of Our Declarative Practices*.
- Wedgwood, R. (2002). The aim of belief. *Philosophical Perspectives*, 16, 267–297.
- Williamson, T. (2000). Knowledge and its Limits. In *Knowledge and its Limits*. Oxford University Press.

CHAPTER 4

Can the Norm of Assertion be Functional?

Several epistemologists hold that there is a constitutive link between the norm of assertion and the speech act of assertion. I present a problem for a novel way of specifying this view, namely etiological functionalism about assertion. Whether the speech act of assertion has the etiological function of generating this or that epistemic good is a contingent matter, determined by the specific selection pressures the practice was selected by. Therefore, etiological functional norms cannot fulfill the constitutive role that epistemologists interested in the norm of assertion have in mind. I offer an alternative path towards a constitutivist functionalism. Functionalists should abandon etiological functions and endorse artifactual functions to account for the constitutive norm of assertion. Such an artifactual functionalist view provides a distinctive path to derive the content norm of assertion, no matter the possible etiological stories of the means by which speakers are able to assert in this or other worlds.

1. Overview

Consider a schema for the norm of assertion:

ϵ -Norm: It is appropriate to assert p only if $\epsilon(p)$

Let ϵ be a placeholder for an epistemically relevant property. Several epistemologists hold that a norm of this kind is fundamental to account for what the speech act of assertion is. Being governed by this norm is essential to the speech act of assertion, such it uniquely individuates assertion from other speech acts. Call this claim constitutivism about assertion.

Constitutivism: There is a norm of assertion ϵ -N such that ϵ -N essentially individuates the speech act of assertion.

Any constitutivist view must meet the metanormative task of explaining what kind of norm is at stake (is it a game-like norm, a social norm, an instrumental norm, etc.?) and the normative task of defining ϵ (is it knowledge, justification, truth, etc.?).¹

Here I evaluate the *metanormative* prospects of a novel view on the norm of assertion, namely etiological functionalism.² According to this view, the normativity of assertion is explained by the fact that assertion exists because ancestor token assertions were selected due to their epistemically beneficial properties. Hence, assertion developed to fulfill the etiological

¹ See Fix (2021), Katsafanas (2018), Korsgaard (2009), and Tubert (2010) for constitutivist examples in other metanormative domains. Under this general parsing, the classic Williamsonian view that assertion is defined by game-like norms is one possible constitutivism (cf. Williamson (2000), Stanley (2008), McKinnon, (2015), Goldberg, (2015)). There can be other constitutivist views appealing to other kinds of norms.

² I focus on Simion and Kelp's work (2021; 2020). See Graham (2020) and Millikan (1987) for similar arguments. Insofar they aim to defend a constitutivist functionalism about assertion, my arguments apply to their view as well.

function of generating the epistemic good it was selected for; and assertions that fail to generate this good fail to fulfill their function. This functional norm, the argument goes, comes metaphysically first than, and has explanatory priority over, other norms that may also govern assertion.

One may ask, however, can this functional norm be *constitutive* of assertion? I present the *contingency problem* for those who want to answer positively (§3). Whether assertion was etiologically selected to generate knowledge, true belief, or any other epistemic good is a matter of how specific selection pressures interacted with the practice of assertion and the environments in which it was selected. Hence, etiological functionalism can only provide a contingent link between the speech act of assertion and a functional norm. It will thus fail to substantiate **Constitutivism** above. If there is anything to the project of determining an essentially constitutive link between assertion and the unique norm that governs it, etiological functionalism will not do.

Of course, functionalists can capitulate and renounce the essentialist claim of **Constitutivism**. Perhaps functionalism should only be interested in the norms that govern assertion in *this* world. As I argue in (§4), however, this move underestimates the contingency problem. It renders functionalism as a project quite different from what philosophers take themselves as doing when discussing the norm of assertion.

Not all is lost for a constitutivist functionalism, though. I argue that the functionalist should renounce the *etiological* component of her project (§5). The constitutive normativity of assertion stems from an *artifactual* function determined by the aim for which a token utterance is intentionally (instead of historically) selected in a purposeful conversation. If we think of assertions as artifacts of this sort, then we can be functionalists and constitutivists: the functional norm that governs token assertions is explained by the instrumental purpose that constitutively defines the act-type of assertion.

2. A Primer on E-Functionalism

Let's begin with the idea of etiological functions. I take the following definition from Simion & Kelp (2021: 74):

E-Function:

A token of type T has the e-function of producing effect E in system S iff

(EF1) Past tokens of T produced E in S's ancestors

(EF2) Producing E benefitted S's ancestors

(EF3) Producing E's having benefitted S's ancestors contributes to the explanation of why T exists in S.³

Etiological functions are most successful when explaining the teleological properties of biological traits. Why does your heart pump blood? Because it is its function. And why is that its function? Because having hearts that pumped blood was beneficial to your ancestors, and this benefit helped them survive. In turn, their survival plus reproduction passed down genes responsible for blood-pumping hearts. Hence, the fact that hearts benefitted your ancestors, plus natural selection, explain the fact that your heart pumps blood.

The beauty of etiological functionalism is that it explains normativity in a perfectly natural world. The fact that blood-pumping hearts were naturally selected because of the benefits they yielded is compatible with the causal order of the world. It also tells us what hearts *should* do. They should pump blood. Even hearts that do not pump blood have the function of doing so. They are just not good hearts. Hence, when it comes to biological traits, e-functions yield *e-functional norms*, grounded on genetic inheritance and natural selection.

When it comes to assertion, the functionalist wager is that, as Kelp and Simion put it, e-functions yield an “independently motivated account of the normativity of assertion” (2021: 97). The practice of assertion exists because past token assertions were etiologically selected due to their epistemically beneficial properties. Therefore, the norm of assertion is best understood as an e-functional norm defined by the epistemic benefit assertion was etiologically selected to generate.⁴

Specifically, Simion and Kelp hold that the practice of assertion was selected because it generates knowledge and, therefore, the knowledge norm of assertion is “etiologically constitutively associated with the speech act of assertion, in virtue of its function of generating knowledge in hearers” (2020: 60).⁵ Yet, since I focus on functionalism as a metanormative story, let us call e-functionalism any view holding that the norm of assertion is an e-functional norm, no matter whether it is defined by knowledge or other epistemic benefit.

E-functionalism offers a coherent metanormative story about the norm of assertion on three accounts. First, it can answer Maitra's (2011) *Hard Question*: ‘What kind of failure

³ The issues raised here should apply to similar views when applied to assertion (e.g., Godfrey-Smith, 1994; Graham, 2014; Millikan, 1987; Neander, 1991).

⁴ Note that there are different plausible ways in which a trait may be etiologically selected. Classically understood, etiological selection happens over generations. Yet, there can be ‘short-story’ etiological that appeals to, for instance, learning processes, maintenance selection, or even by metabolic activity. What matters is that there is a positive feedback loop for which that non-intentionally reinforces the presence of the trait at stake (cf. Simion 2021; Graham 2014; Buller 2008; Macdonald & Papineau 2006). Since, these are still contingent processes, and thus are subject to my argument below, I bracket this wrinkle and refer to e-functionalism as a general story.

⁵ Although see §4 for a crucial caveat on what they mean by constitutivism.

constitutes an assertion that does not comply with its norm?’ Assertions that do not bring about this epistemic good fail to fulfill their e-function. Second, it explains why bad assertions are still assertions. Just as dysfunctional hearts are still hearts, assertions that don’t succeed in generating the epistemic good for which they were selected are still assertions. And third, it can explain why sometimes inappropriate assertions can be better off than others. A heart that pumps orange juice is not fulfilling its function. Yet, it is functioning normally. Similarly, some assertions may not fulfill their function because they were issued outside the normal conditions in which they would generate the epistemic good for which the general practice of assertion was selected.⁶

3. The Contingency Problem

Recall the question at stake: can e-functional norms be constitutive of assertion? To introduce the contingency problem for a positive answer to this question, consider two functionalist options.

First, Millikan (1987) and Graham (2020) argue for the truth norm. Their argument is simple. Generating true beliefs through assertions presumably benefits a communicative system and it is plausible that this is the benefit assertion was selected for. Second, Simion & Kelp (2021) offer a similar account for the knowledge norm. Knowledge is valuable and thus beneficial.⁷ Hence, it is plausible that assertion was selected because it generated knowledge. Further, in discussing the truth norm, they hold that since knowledge is a better good than truth, it is more plausible that assertion was selected because it generated knowledge, not true beliefs.

Note, however, that natural or cultural selection does not entail that the best or better benefit must be selected. What benefit is selected is determined by the actual properties of the ancestor tokens and their interaction with a given environment. It would be better for us if our lungs were selected to absorb oxygen inside and outside of water. But they were not selected to do this because the genotypical and phenotypical properties of our ancestors and their interaction with the environment precluded that possibility. So, even if knowledge is better than true belief, it may be that the fact that past assertions generated true belief is enough to ensure the functional norm of assertion.⁸ That sometimes assertion generates

⁶ Critics of functionalism in the social sciences argue that it is not clear that social practices are always explained by selection mechanisms parallel to natural selection (Elster, 2007; Giddens, 1986). Assertions don’t have genes and there are several, non-benefit-grounded ways in which communities select and maintain their social practices. Simion and Kelp are notably silent about this aspect of their view. As Graham, (2020) suggests, however, cultural evolution and social norms can ground the e-functional norms that govern assertion. I assume that such a story can be told. Although see Bicchieri (2017) for examples of social norms maintaining practices that do not benefit—and even hurt—a community.

⁷ Though see (Papineau, 2021).

⁸ Here’s a plausible story. Presumably, given natural instances of Gettierized true belief, the number of true assertions is larger than the number of knowledgeable assertions. This may have eased the selection of true assertions rather than knowledgeable

knowledge may be an added benefit, but not the benefit assertion was selected for—just as the fact that my lungs allow me to meditate doesn't entail that they were selected for this.

Of course, this argument needs more to be substantiated. What matters for my purposes, however, is that it is possible that assertion was selected for generating true belief and it is possible that it was selected for generating knowledge. Crucially, this means that whether the functional norm of assertion is truth or knowledge is a contingent matter. There are possible (plausible, or even overwhelmingly plausible) worlds in which assertion was etiologically selected to generate either *and not the other* epistemic good—or even to generate other goods like certain, justified, or rational belief.

More radically, in admittedly distant worlds, assertion may have been selected for non-epistemic benefits. In some words, ancestor assertion tokens may have been selected because they benefitted, or avoided punishment from, a powerful subset of the community; or because they were attractive to another subset of the population. In those cases, the norm of assertion would be 'say congenial things' or 'say attractive things.'

Recall, however, the starting point of the constitutivist about assertion. Constitutivists interested in the norm of assertion hold that it is essential to the speech act of assertion in a way that uniquely individuates it. Put otherwise, philosophers interested in the norm of assertion begin with the notion that assertion essentially involves being governed by said norm. Nevertheless, the contingent relationship that the e-functionalism can derive between assertion and its e-functional norm will not fulfill this conceptual role.

Assertion could be selected because it yielded different benefits and thus it could be governed by different e-functional norms. If all there is to the norm of assertion is etiological selection, then it is contingent to the act of assertion. This means that e-functional norms do not constitutively what assertion is. Assertion cannot be the act that is essentially individuated by its being governed by an e-functional norm. Hence, the link between assertion and its e-functional norm(s) is not the link that the constitutivist is looking for.⁹

4. Renouncing Essentialism

Moved by these considerations, Kelp and Simion describe their view as a "light" constitutivism. They note that all they mean by saying that the e-functional norm of assertion is constitutive of assertion is that the e-function of generating knowledge plays, as a matter of

assertions, given that that, arguably, whether a communicated true proposition is useful or not doesn't depend on how the speaker got to believe it but whether the audience's purposes are aided by it being true.

⁹ Could functionalism allow an essential link between assertion and some or other epistemic good? Note two considerations against this suggestion. First, this is a far cry from the initial norm-centered project. Defenders of the K-norm and the T-norm say it is theirs and no other norm that distinguishes assertion. Second, recall that in some world's assertions may be selected for non-epistemic benefits. Those worlds preclude e-functions to essentially link assertion with an epistemic good.

fact, “a vital role” in explaining why assertion continues to exist. Hence, “it may be that assertion has the function of generating knowledge contingently” (2019: 69).

Say, then, that the e-functionalist accepts something like the following claim.

Non-Essential Constitutivism: There is a norm of assertion ϵ -N such that ϵ -N individuates the speech act of assertion.

Without saying more, however, it is not exactly clear why this amounts to a constitutivist view.¹⁰ Still, independently of whether it deserves the ‘constitutivist’ label or not, determining whether an etiological function has a vital role in maintaining—without essentially defining—a practice is a perfectly legitimate project. Perhaps functionalists about assertion shouldn’t be interested in accounting for an essential relation between the norm of assertion and the speech act of assertion. Perhaps functionalists should be content with etiological explanations of why assertion has the function of offering a single epistemic good in *this* world.

Note, however, that this project is quite different from what philosophers interested in the norm of assertion have in mind. For starters, disputes on the norm of assertion are framed as disputes on the essential norm that *uniquely* characterizes the speech act of assertion. Yet, under the non-essentialist rendering of e-functionalism, disputes on the norm of assertion amount to empirical disputes on whether the specific interaction of ancestor assertion tokens with the world happened to favor one or another epistemic good. And just as determining whether marriage has the e-function of securing property inheritance or ensuring the means of childbearing, adjudicating the e-function of assertion is a project contingent on the specific histories of the communities that enact this practice.

Crucially, this means that e-functionalists must be open to the fact that assertion may have been culturally selected for different purposes in different communities. Hence, evidence that in a given community assertion has the function of generating a given epistemic good will not warrant that assertion has the same function across communities. This is even more pressing once we note that functionalism must also be open to the fact that social practices can evolve to fulfill different functions at different times—just as wings have different functions for chickens than for ravens. Even if assertion emerged as a mechanism to fulfill the function of generating a single epistemic benefit, this practice could have evolved to fulfill different functions as communities developed and diverged.

Therefore, if all there is to the norm of assertion is etiological selection, it is not so clear that functionalism should aim to individuate a *unique* norm of assertion, even in this world. It seems plausible that, as Cappelen (2011) suggests, there are a wide number of assertoric practices, none of which has the exclusive privilege of being called *the* act of assertion. If so,

¹⁰ It is certainly incompatible with metanormative constitutivism as exemplified by those cited in n.1.

whether one assertoric practice or another deserves the label of assertion comes down to a terminological debate, not to an insight on the epistemic nature of our communicative practices. Alternatively, to show knowledge—or any other epistemic good, for that matter—is the unique norm of assertion, it must be shown that a) assertion is the kind of practice selected by cultural selection to provide a singular epistemic good; b) it is the unique such practice; and c) its uniqueness is maintained across cultures and through time. This remains to be done.¹¹

Simion and Kelp appeal to their light constitutivism to develop a knowledge-first account of assertion that captures the main intuitions of essentialist constitutivism. Yet, we have seen, under etiological functionalism, knowledge is not constitutive of assertion across worlds, nor it is obvious that knowledge individuates *the* act of assertion in this world. Even if there is indeed an e-functional link between actual assertions and knowledge, it seems to be a rather weak one.

Here is another way of viewing the problem. Constitutivism about assertion is interested in offering an account of what assertion *is*. However, etiological selection can only offer contingent stories about how specific systems develop certain traits. Since these stories can and do vary widely depending on the specific relations of a given system with its (everchanging) environment, it is implausible that they will offer stable and unified accounts of what the distinct speech act of assertion amounts to.¹² In the next section, I offer an alternative path to determine what assertion is. Still, we will see, there is plenty that etiological functionalism can offer—just not an account of the nature of assertion.

5. Renouncing Etiology

Etiological functions, we have seen, are quite successful in dissipating the mystery of teleological properties seemingly designed to fulfill a purpose in a world without an agency that does so. Yet, as Neander (1991) suggests, there is another kind of teleological feature for which there is no such mystery: the functionality of intentionally made artifacts.¹³ Artifacts are indeed designed by an agency—*our agency*.

¹¹ Note that while this kind of evidence is not clearly available for assertion, it is for hearts. This is because, as Marr points out, functional analyses require a prior description of the mechanism's defining properties. While nobody disputes whether hearts are meant to pump blood or not, there is certainly disputes about whether assertion is meant to generate knowledge or not. As I argue below, it is important not to confuse an analysis of the defining properties of a mechanism with its functional analysis.

¹² This charge is similar (but not the same) to causal-constitutive problems in philosophy of mind, cf. Churchland (1994), Kirchhoff (2015). Etiological functionalism can (and certainly does) offer informative causal explanations of how it came to be that assertions are possible in our world. It will yield constitutive explanations what the act of assertion amounts to.

¹³ Neander notes that she diverges from other classic functionalists on this issue (1991: 462).

If I select an object to achieve a given purpose, I imbue it with the function of fulfilling that purpose. A rock can have the function of hammering a nail insofar as I have selected for this purpose. While the blood-pumping function of a heart is grounded on natural selection, the hammering function of that rock is grounded on my having intentionally selected it to perform the act of hammering. The first ever object intentionally selected to perform the act of hammering had the function of hammering. No ancestor needed. And, crucially, the act of hammering amounts to the same act across worlds, no matter the selection stories that led to whoever hammers being able to hammer.

Likewise, my suggestion here is that assertion can be essentially individuated by the purpose for which an utterance is intentionally selected in conversation. This approach allows a distinctive way of deriving an epistemic norm of assertion in a way that is not contingent on the possible histories by which it came to be possible to assert in this or any other world.

Begin with a plausible observation about assertion, compatible with a wide slew of treatments of assertion across disciplines. Assertions occur whenever a speaker presents a given utterance as communicating information that advances the purposes of the conversation. What distinguishes assertions from sarcasm, irony, or even sleep talking is that an utterance is rationalized as having the role of actively communicating a piece of information, given the speaker's and audience's acknowledgement that they are engaged in purposeful conversation.¹⁴

This, I submit, should be the starting point of an artifactual account of assertion.

Artifactual Assertion: A token utterance asserts *p* whenever it is presented as aimed at communicating the information that *p* to advance the purposes of the conversation.

To assert, a speaker issues an utterance that, given the purposes of the conversation in which it is embedded, is interpreted as aimed at communicating information.¹⁵ If you and I are cooking pasta and looking for salt, uttering “it’s in the cupboard” counts as asserting that the salt is in the cupboard because that is the warranted interpretation of my act, given the purpose of our conversation.¹⁶

¹⁴ Of course, what this ultimate purpose is can vary widely. It can be communicative (sharing information about where Neil Young was born) or non-communicative (getting to the movies); it may even be a vague purpose (having a pleasant evening conversation). Note too that even contentious conversations have shared purposes, even if they are imposed by one party you may be in a conversation with me because you are afraid of what would happen otherwise. Likewise, two dueling lords may want to kill each other but still be motivated to engage in conversation find a place to duel. What matters is that an utterance is rationalized as aimed to provide information given the mutually acknowledged aims of the conversation.

¹⁵ What about insinuations? Insinuations exploit the vague practical role of an utterance in the purposeful structure of the conversation. While it could be that they are communicating *p* it could also be that they are communicating *q* (such that not-*p*). While assertions require a clear practical role of an utterance, insinuations preclude it. See (Pinker, Camp, Vesga).

¹⁶ What about diary entries, soliloquies, or even on off assertions.

The aim of communicating information, in turn, imbues token assertions with the function—and its corresponding normativity—of providing information that advances the purposes of the conversation. A successful assertion indeed provides the information it was presented as communicating. Only if the salt is in the cupboard does my assertion fulfill the aim for which I have presented it in conversation. Yet, assertions may be unsuccessful. Perhaps the speaker is deceptive or incompetent and failed to indeed offer the information they presented themselves as communicating. If the salt is not in the cupboard, my assertion fails to achieve what I have presented it as achieving.

Importantly, note that the artifactual account of assertion can say more about the epistemic profile of the aim of assertion than just a truth condition. The key strategy is to characterize assertion as an act individuated by its aim: that of providing information. This renders the normativity to which a token assertion is subject in terms of success. An assertion can be normatively assessed on whether it successfully achieved the aim it was presented as having. And, crucially, since the act of providing information is an epistemically significant act (for it involves representations of how the world is), the practical profile of successful action will determine the epistemic profile of the special case of assertion.

Consider the following case for a derivation of the epistemic profile of the norm of assertion. We have seen how the practical profile of successful assertion requires truth. If one successfully ϕ -s one actually ϕ -s. Just as I only successfully hammer a nail if I actually hammer a nail, I only provide the information that the salt is in the cupboard if the salt is in the cupboard.

There's more. One successfully ϕ -s only if one intentionally ϕ -s, and intentional action requires a sufficient degree of belief. At the very least, one must not believe that one will not ϕ .¹⁷ Therefore, to successfully assert one must have a sufficient degree of belief that one will indeed provide truthful information—which requires sufficient credence of what one asserts.

Further, one successfully ϕ -s only if there is a proper connection between one's agency and one's ϕ -ing. I have not successfully hammered if I suddenly wake up and realize that I was hammering in my sleep, nor if I drop my hammer and it happens to perfectly land on a nail.¹⁸ Similarly, there must be a proper link between my *epistemic* agency (my reasons or my evidence, for instance) and the fact that I am asserting. This link arguably requires justification.

And, finally, this link between one's agency and one's action requires modal robustness. Just as I cannot intentionally win the lottery, it is odd to say that I successfully hammered a nail if I closed my eyes and gave one random guess of where the nail was. Even if the nail ends up being nailed, my action could have gone awry too easily; I lacked sufficient

¹⁷ Cf. Bratman (1987), Broome (2013), Pavese (2017). But see Nuñez (2020) for complications.

¹⁸ Cf. Bratman (1987), Davidson (2001), Mitchell (1982), Tännsjö (2009).

control over my action.¹⁹ Similarly, the link between my epistemic agency and my successful assertions must ensure that my assertions are successful (and thus true) in a sufficient number of similar worlds—that is, my assertions should be safe.

So, insofar as successfully φ -ing requires actually φ -ing, sufficient belief that one will φ , a proper link between one's agency and one's φ -ing, and this link to be sufficiently robust, assertion arguably requires truth, sufficient belief, justification, and safety. When individuated by the aim of providing information that advances the purposes of the conversation, assertions may fail in any or all of these dimensions.

Of course, you may disagree. Perhaps successful action requires more or less than what I just suggested. Perhaps these requirements have different consequences for the epistemic profile of assertion. Since my focus here is on the *metanormative* status of artifactual functionalism, I bracket these *normative* concerns about the content of the norm of assertion here. What matters for my purposes is that our disagreement is substantial. It concerns the nature of successful action and the consequences it has for the special case of the epistemically significant act of assertion. Artifactual functionalism offers a stable path to derive the norm of assertion by appealing to what it takes to intentionally act.

Granted, we may develop objects or practices that are better suited to fulfill an artifactual function. We progressively developed clawhammers, mallets, and sledgehammers to better hammer. We progressively developed English, Finnish, and Esperanto to better assert. By the same token, we evolved particular cognitive and physiological means that enable us to hammer and assert. But it is important to distinguish accounts of how specific hammers came to be and how we came to be the kind of animal that is able to hammer from an account of what it is to hammer. It is equally important to distinguish an account of the specific stories of how we came to be able to assert from an account of what it is to assert.

This means that there is plenty of work for etiological functionalism in forming a full account of human assertion, for it is primed to offer explanations that meet the former task. It can offer accounts of the different cognitive and physiological features that enable us to assert. It will also offer accounts of how and why communities developed practices that govern, shape, and constrain our assertions. These accounts will yield the standards by which tokens of these features and practices are better or worse performing the functions that enable us to assert. Again, however, these stories will not tell us what the act of assertion amounts to. They will not yield a constitutivist picture of assertion.

In contrast, artifactual functionalism brings the act of assertion and its epistemic requirements much closer together than etiological functionalism can ever do. The fact that assertion is governed by certain epistemic norm owes to the very nature of purposeful action,

¹⁹ Cf. to Mele (1995) and Mele & Moser (1994), Shepherd (2021), Greco (2016).

and not to the contingent interactions that happened to ensue between asserters and their environment. Individuated by their aims, ‘hammering’ and ‘asserting’ are rigid designators of the same act-types, no matter how they are performed, the means they are performed with, or the specific histories that enabled agents to perform them.²⁰ Crucially, this means that an artifactual account of assertion does not face the contingency problem. The act of assertion is *constitutively* defined by its purpose, no matter the world it is performed in. And whether a token assertion fulfills its artifactual function or not comes down to whether it successfully achieved the purpose that characterizes the act of assertion.

Keep in mind that this is only the metanormative side of the project. The fact that assertion is an act-type governed by an instrumental norm has explanatory priority in explaining the functional norm that governs token assertions. Put otherwise, the core normativity of assertion is instrumental. It is still up for debate (and development) whether the normative story I have offered here holds.²¹ Perhaps this profile amounts or falls short from knowledge. Perhaps the required epistemic profile will be contextually determined by the practical stakes of the conversation.

Crucially, however, this dispute is not a terminological debate on what practice should be labeled assertion or not. Instead, it is a dispute on the nature of the act of assertion and the purpose that individuates it from other actions, no matter the world it is performed in.²² In other words, it is a dispute about what assertion constitutively requires from an agent that performs it.

References

- Bicchieri, C. (2017). *Norms in the wild: How to diagnose, measure, and change social norms*. Oxford University Press.
- Bratman, M. (1987). *Intention, Plans, and Practical Reason* (Issue 2, pp. 198–199). Cambridge, MA: Harvard University Press.
- Broome, J. (2013). *Rationality Through Reasoning*. Wiley-Blackwell.
- Buller, D. J. (1998). Etiological Theories of Function: A Geographical Survey. *Biology and Philosophy*, 13(4), 505–527.

²⁰ There certainly is a specific story of how we evolved as a species to be able to perform the act of assertion. This story will not yield the constitutive norm of assertion, however.

²¹ Once the normativity of assertion is seen as mainly instrumental, the constitutive link between successful action and epistemic control comes to the fore in the project of determining the epistemic requirements of successful assertion (cf. Williamson 2000; Gibbons 2001; Greco 2016; Pavese 2018). The specifics of this normative project, however, exceed the scope of the metanormative focus I have adopted here.

²² Note that defining the act of assertion in this way does not require appealing to a game-like practice, as it is assumed by Williamsonian constitutivists.

- Cappelen, H. (2011). Against Assertion. In J. Brown & H. Cappelen (Eds.), *Assertion: New Philosophical Essays*. Oxford University Press.
- Churchland, P. S. (1994). Can Neurobiology Teach Us Anything about Consciousness? *Proceedings and Addresses of the American Philosophical Association*, 67(4), 23–40.
- Cohen, L. J. (1982). What Is Necessary for Testimonial Corroboration? *The British Journal for the Philosophy of Science*, 33(2), 161–164. JSTOR.
- Davidson, D. (2001). *Essays on Actions and Events*. Oxford University Press.
- Elster, J. (2007). *Explaining Social Behavior: More Nuts and Bolts for the Social Sciences*. Cambridge University Press.
- Fix, J. D. (2021). Two Sorts of Constitutivism. *Analytic Philosophy*, 62(1), 1–20.
- Giddens, A. (1986). *The Constitution of Society: Outline of the Theory of Structuration* (1st edition). Polity.
- Godfrey-Smith, P. (1994). A modern history theory of functions. *Noûs*, 28(3), 344–362.
- Graham, P. J. (2014). Functions, Warrant, History. In A. Fairweather & O. Flanagan (Eds.), *Naturalizing Epistemic Virtue* (pp. 15–35). Cambridge University Press.
- Graham, P. J. (2020). The Function of Assertion and Social Norms. In S. Goldberg (Ed.), *The Oxford Handbook of Assertion* (pp. 725–748). Oxford University Press.
- Graham, P. J. (2014). Functions, Warrant, History. In A. Fairweather & O. Flanagan (Eds.), *Naturalizing Epistemic Virtue* (pp. 15–35). Cambridge University Press.
- Greco, D. (2016). Safety, Explanation, Iteration. *Philosophical Issues*, 26(1), 187–208.
- Goldberg, S. C. (2015). *Assertion: On the Philosophical Significance of Assertoric Speech*. Oxford University Press.
- Graham, P. J. (2020). The Function of Assertion and Social Norms. In S. Goldberg (Ed.), *The Oxford Handbook of Assertion* (pp. 725–748). Oxford University Press.
- Hawthorne, J., & Stanley, J. (2008). Knowledge and Action. *Journal of Philosophy*, 105(10), 571–590.
- Katsafanas, P. (2018). Constitutivism About Practical Reasons. In D. Star (Ed.), *The Oxford Handbook of Reasons and Normativity* (pp. 367–394). Oxford: Oxford University Press.
- Kelp, C., & Simion, M. (2021). *Sharing Knowledge: A Functionalist Account of Assertion*. Cambridge University Press.
- Kirchhoff, M. D. (2015). Extended Cognition & the Causal-Constitutive Fallacy: In Search for a Diachronic and Dynamical Conception of Constitution. *Philosophy and Phenomenological Research*, 90(2), 320–360.
- Korsgaard, C. M. (2009). *Self-Constitution: Agency, Identity, and Integrity*. Oxford University Press.
- Maitra, I. (2011). Assertion, Norms, and Games. In J. Brown & H. Cappelen (Eds.), *Assertion: New Philosophical Essays* (pp. 277–296). Oxford University Press.
- Mele, A. R. (1995). *Autonomous Agents: From Self Control to Autonomy*. Oxford University Press.

- Mele, A. R., & Moser, P. K. (1994). Intentional Action. *Noûs*, 28(1), 39–68.
- McKinnon, R. (2015). *The Norms of Assertion: Truth, Lies, and Warrant* (2015 edition). Palgrave Macmillan.
- Millikan, R. G. (1987). *Language, Thought, and Other Biological Categories: New Foundations for Realism* (Reprint edition). MIT Press.
- Mitchell, D. (1982). Deviant Causal Chains. *American Philosophical Quarterly*, 19(4), 351–353.
- Nagel, E. (1961). *The Structure of Science: Problems in the Logic of Scientific Explanation*. New York, NY, USA: Harcourt, Brace & World.
- Neander, K. (1991). The teleological notion of “function.” *Australasian Journal of Philosophy*, 69(4), 454–468.
- Núñez, C. (2020). Requirements of intention in light of belief. *Philosophical Studies*, 177(9), 2471–2492.
- Papineau, D. (2021). The disvalue of knowledge. *Synthese*, 198(6), 5311–5332.
- Pavese, C. (2018). Know-how, action, and luck. *Synthese*.
- Simion, M., & Kelp, C. (2020). Assertion: The Constitutive Norms View. In S. Goldberg (Ed.), *The Oxford Handbook of Assertion*. Oxford University Press.
- Tännsjö, T. (2009). On deviant causal chains: no need for a general criterion. *Analysis*, 69(3), 469–473.
- Stanley, J. (2008). Knowledge and Certainty. *Philosophical Issues*, 18(1), 35–57.
- Tubert, A. (2010). Constitutive Arguments. *Philosophy Compass*, 5(8), 656–666.
- Williamson, T. (2000). Knowledge and its Limits. In *Knowledge and its Limits*. Oxford University Press.

CHAPTER 5

Talking With Others: The Cooperative Norm of Assertion

Abstract: Constitutivism about assertion holds that there is a constitutive link between the norm of assertion and the very nature of the speech act of assertion. This family of views has been mostly concerned with what I call here the normative question about assertion: should speakers, for instance, know what they assert, should they have justified belief, or should they say true things? Here I suggest a novel path to answer this question by way of addressing another significantly less visited question: what kind of norm is the norm of assertion? Call this the metanormative question about assertion. I argue that the constitutive norm of assertion is not game-like, social, conventional, or functional. Rather, it is an instrumental norm, stemming from the normativity of cooperative action. This view neatly accommodates the non-constitutive role of other communal norms in the practice of assertion. In turn, this cooperativist view provides a distinctive path to specify the epistemic profile of the norm of assertion. The link between intentional action and success suggests that the norm of assertion requires asserted contents to be believed, true, justified, and safe. Whether this profile amounts to knowledge, however, will depend on your epistemological allegiances.

1. Overview

Consider a paradigmatic assertion. You and I are deciding where to eat, and I utter “There’s a restaurant on fifth and Cascadilla.” Say, however, that there is no restaurant there—turns out I was merely guessing. Clearly, there is something inappropriate about my act. You are entitled to criticize it because I lacked adequate epistemic standing with respect to what I said. This strongly suggests that there is an epistemically significant norm of assertion that adjudicates between proper and improper assertions.

A venerable tradition in epistemology holds that there is a constitutive link between such a norm of assertion and the very nature of the speech act of assertion. Being subject to this norm is essential to the speech act of assertion in a way that uniquely individuates it from other speech acts. Call this general claim constitutivism about assertion.

Here is a sample schema of this norm:

ϵ -norm: It is appropriate to assert p only when $\epsilon(p)$

Let ϵ be an epistemically relevant property. Most constitutivists about assertion focus on the *normative* task of specifying this property—is it, for instance, knowledge, truth, or justified belief? Yet, not enough attention has been paid to the *metanormative* question of specifying the kind of norm at stake: what in the world makes it the case that assertion is constitutively subject to this norm? My aim here is to offer a distinctive pathway to answer the normative question about the content of the norm of assertion by way of providing a novel constitutivist account of its metanormative status.

I begin by motivating a metanormative shift from extant accounts in the literature (§2). Most constitutivist views fall into what I call Communal Constitutivism—the claim that the constitutive norm of assertion is a communal norm generated by the collective interactions of the members of a community. Examples of this kind of norms are game-like norms, conventions, social norms, and etiological functions. Nevertheless, as I argue, no communal norm can be the constitutive norm of assertion. If a norm n of assertion is indeed constitutive of assertion, it must be possible to explain why n , and not another norm m or l , exclusively and constitutively governs the unique act of assertion. Yet, no communal norm has this kind of authority over other possible communal norms.

There is an alternative story that can fulfill the metanormative role constitutivists are interested in. Inspired by another important tradition that gives center stage to non-language-specific practical rationality in explaining communication, I offer a cooperativist account of assertion. To assert, I contend, is to issue a signal that, given the shared goals that structure the conversation in which it is embedded, must be interpreted as cooperatively offering information as means to achieve said goals (§3). Under this cooperativist picture, the norm of assertion is an instrumental norm derived from the distinctive instrumental rationality of cooperative action. In this sense, the norm of assertion follows from the constitutive aim of assertion as a cooperative act (§4).

This view not only avoids the adjudication problem faced by other norms; it also explains the non-constitutive role of conventional, etiological, social, and game-like norms with respect to the act of assertion. These norms guide, specify, motivate, and constrain—but do not constitute—the cooperative interactions that enable the act of assertion (§5). Hence, the cooperativist view has the further metanormative advantage of yielding a unified picture of the various constitutive and non-constitutive norms of assertion.

Finally, I explain how the cooperativist view can also fulfill the normative task of stating the content of the norm of assertion (§6). The link between action—cooperative action, in this case—and success entails that the instrumental standard of cooperatively φ -ing requires i) actually φ -ing, ii) a belief constraint concerning whether one will φ , iii) a proper connection between the agent and her φ -ing, and iv) that this connection is modally robust. This, I argue, suggests that, when it comes to the action of cooperatively offering information (i.e., asserting), the instrumental standard of assertion requires truth, belief, justification, and safety. Arguably, these conditions are at least entailed by knowledge. Whether they amount to knowledge, however, will depend on the reader's epistemological allegiances. I conclude by parsing out some paths one could take with respect to the claim that knowledge is the norm of assertion.

Before I begin, a clarificatory note. My objective here is to defend a constitutivist picture of assertion immune to the challenge I pose and to other challenges issued against classic

constitutivist views of assertion. So, my claim is that, *if* you are going to be a constitutivist about assertion, you are better off accepting my cooperativist view.¹ Given the sheer popularity of constitutivist views of assertion, I take it that this project ought to be of interest to an important number of epistemologists. Yet, I postpone the task of defending this view against other non-constitutivist views of assertion in the literature.²

2. Communalism and its Discontents

By ‘norm’ let us understand a standard against which a token action is deemed appropriate or inappropriate.³ Hence, in its most general, constitutivism about assertion can be parsed as follows:

Constitutivism: There is a norm of assertion ϵ -N, such that ϵ -N essentially individuates the unique act of assertion.

As I have said, a vital task of any constitutivist account of assertion is to say what kind of norm or standard is at stake; what in the world makes it the case that this standard is in place? Virtually all constitutive accounts of assertion in the literature fall into what I call here communal constitutivism, which is characterized by the following additional claim:

Communalism: ϵ -N is a communal norm.

In a gist, a communal norm is a standard whose existence is explained by the collective behavior of a given community. This will be clarified by example.

Most influentially, Williamson (1996, 2000) develops a constitutivist account of assertion by appealing to the Austinian/Searlean idea that some socially established and enacted rules of conduct give rise to distinctive social practices.⁴ An example of this kind of rule is the rule in chess stating that switching your pawn for a queen when the former reaches the far side of the board counts as promoting said pawn. Once this rule is accepted by chess players, it determines whether a given move counts as a felicitous (or infelicitous) instance of a distinctive act with a determined set of consequences: pawn promotion. Other examples of game-like norms are the norms that govern the acts of making someone a priest, declaring a

¹ See Goldberg (2019) for a thorough defense of what I am calling here constitutivism against non-constitutive views of assertion.

² Examples of non-constitutive views of the norm of assertion are views that hold that assertion is an expression of an attitude (Green 2010; Bar-On 2007), that there is more than one norm of assertion (Cappelen 2011), social-consequences views of assertion (Brandom 1983; Lance & Kukla 2009), and that assertion is not characterized by a norm (Paguin 2004, 2008).

³ Henceforth, I use ‘rule,’ ‘norm,’ and ‘requirement’ interchangeably. There are important differences between these three notions. But at their core, they all mark a standard of appropriateness.

⁴ As Searle (1969) has it, assertion is a social practice constituted by rules that “create or define new types of behavior” (p. 33) by defining that “X counts as Y in context C” (p. 35).

holiday, or hitting a homerun. That assertion is just one more of these practices is the typical metanormative story in the constitutivist literature.⁵

Nevertheless, there are other kinds of norms to which the communal constitutivist can appeal. For instance, one may think that the constitutive norm of assertion is a convention. That is, one may think that the norm of assertion owes its existence to a communal pattern of behavior that guides and streamlines coordination between members of a community, like the convention of driving on the right-hand side of the road, or the convention of taking the garbage out on Thursdays (Lewis 1980; Geurts 2018; Marmor 2009).⁶ Or one may also think that the constitutive norm of assertion is a social norm—a standard maintained by informal sanctions that motivate coordination in agents who would not coordinate otherwise, just as the norm of queuing at the luncheon, or of wearing a tie to a wedding (Bicchieri 2006, 2016; Brennan et al. 2016; Martinez ms).⁷

And, as has been more recently argued, one may think the norm of assertion is an etiological functional norm, or e-functional norm for short (Kelp & Simion 2021; Graham 2020). The persistence of some social practice *s* in a community may be explained by the fact that *s* was historically selected because it typically yields benefit *b* (Millikan 1987, Pettit 1996). If so, *s* has the e-function of generating *b*. In some communities the practice of marriage may fulfill the e-function of ensuring property inheritance; in some other communities, myths about food have the function of avoiding poisonous preparations (Henrich 2016). In turn, one can judge whether token actions of an e-functional social practice are appropriate by evaluating whether they indeed fulfill its e-function. Hence e-functions generate e-functional norms. So, one may think that the constitutive norm of assertion stems from assertion's e-function of yielding a given epistemic good, given processes of selection that gave rise to the practice of assertion.⁸

As we will see in §5, there is something right about the intuition that all of these communal (game-like, conventional, social, and e-functional) norms have important normative roles with respect to the act of assertion. However, my argument here is that no communal norm can yield a *constitutive* account of essentially individuating the unique act of

⁵ Now-classic examples of 'game constitutivism' are DeRose (2002), Hawthorne (2004), Weiner (2005), Douvan (2006), Lackey (2007), and Kvanvig (2009). Even arguments against constitutivism of assertion in general have focused on the problems that follow from characterizing assertion as a game-like practice (cf. Maitra 2011; Cappelen 2011; Pagin 2016). See Marsili (2019) for a critical overview of this literature.

⁶ See Fricker (2006) and García-Carpintero (2019) for plausible constructions of the norm of assertion as conventional.

⁷ See Graham (2018), Starr (2018), and Murray & Starr (2021) for considerations of the crucial role of social norms in the practice of assertion.

⁸ Importantly, e-functionalism can appeal to many selection stories apart from cultural selection, such as learning processes, maintenance selection, or even metabolic activity. Since all of these stories are subject to the challenge I pose below I bracket this wrinkle of e-functionalism.

assertion. Put otherwise, **Constitutivism** and **Communalism** above are incompatible and, hence, communal constitutivism is misguided.

To begin parsing the challenge, note that there is significant disagreement on the content of the norm of assertion. Some constitutivists have argued assertion requires knowledge⁹, others take it to require truth¹⁰, yet others take it to require only justification or reasonableness.¹¹ Say you and I are constitutivists, but I consider that the norm of assertion requires an epistemic good ϵ and you that it requires ϵ^* . Given our commitment to **Constitutivism**, this amounts to a disagreement on what assertion is. And, moreover, **Constitutivism** commits us to think that at least one of us is wrong. After all, constitutivists are committed to it being essential to assertion that it is *uniquely* individuated by the norm of assertion. Yet, it is hard to say how a commitment to **Communalism** allows for this result. That is, it is hard to say what exactly gives any communal norm the authority to constitute the unique act of assertion, and thus adjudicate our disagreement.¹²

Consider game-like rules. When we say that people act wrongly or inappropriately in moving their knight diagonally, or reciting the wrong words when ordaining a minister, we do so against pre-established norms that make up these practices. Crucially, these norms are indexed to specific communities that accept them. And, consequentially, these norms do not govern the actions of agents who do not accept them and thus do not engage with the practice. The community of speed-chess players and the community of 960 chess players judge the appropriateness of their acts against different norms. It would be a moot point if they appealed to the rules of their respective games to criticize each other for playing the wrong game. There is nothing within either practice that adjudicates moves outside it. They are just different games.

Similarly, the game constitutivist who argues that assertion is constituted by the game-like norm requiring an epistemic good ϵ and not ϵ^* may as well just be describing the difference between their practice of assertion and another person's or community's practice of assertion. It is not clear what about any of these norms would have the normative privilege of constituting the unique game-like practice of assertion. If assertion is the practice of subjecting one's sayings to *the* game-like norm of assertion, the onus remains on the game constitutivist

⁹ E.g., Williamson (1996, 2000); DeRose (2002); Hawthorne (2006).

¹⁰ E.g., Weiner (2005), Whiting (2012).

¹¹ E.g., Douven (2006); Lackey (2007); Kvanvig (2009).

¹² This disagreement has been noted by others challenging game constitutivism. Pagin (2016), for instance, notes that, if the game constitutivist allows that these norms are implicitly accepted, game constitutivism entails the theoretically loaded posit of a universally accepted game or practice, constituted by a specific norm that we all know how to follow but are mostly unable to express. Since my argument against these views is more general, I allow that game-like norms can be implicitly accepted.

to explain, without circularity or regress, what determines the overall normative authority of this single norm over other candidate norms.

This issue extends to other kinds of communal norms. When we say that people act wrongly or inappropriately when they do not follow the convention of driving in the right-hand lane in England or the social norm of wearing a tie to a wedding, we do it against standards that are contingent on the practices and histories of specific communities. Nobody expects these norms to hold universally on all relevant act-types. The fact that some communities drive on the right side of the road yields no normative authority to say that other communities who drive on the left side of the road are driving wrongly. The fact that some communities developed the social norms of wearing a tie to a wedding does not make wearing a sherwani to a wedding the wrong kind of act. Likewise, the fact that a given community develops marriage to have the function of ensuring property inheritance does not entail that communities in which marriage has the function of ease childbearing are performing the wrong kind of practice.

In the face of two (or more) practices, guided by conventions, social norms, or e-functions, there is nothing about these communal norms that makes one of these practices, and not the other, the correct practice. Yet, that is exactly what the constitutivist of assertion expects from the norm of assertion; something about this norm makes it *the* norm of assertion in a way that essentially individuates it from other candidate norms, and thus individuates assertion from other candidate speech acts.

The communal constitutivist may be tempted to appeal to *actual* practices of assertion. Perhaps through some social process—be it cultural evolution, social agreement, or even force—the practice of assertion came to be indeed governed by a single communal norm. Perhaps most people actually follow your communal ϵ^* -norm of assertion and not my communal ϵ -norm of assertion. And it is against the communal norm of this accepted-by-most practice of assertion that we can settle disagreements of what the norm of assertion is and thus what the act of assertion ultimately amounts to.¹³

For starters, note that this would be a much weaker claim than originally intended. The task of specifying the content of the norm of assertion would be rendered closer to the empirical work of sociologists and anthropologists than to what the epistemologist usually takes herself as doing. That the norm of assertion demands ϵ and not ϵ^* would be an empirical claim about our actual practices. It would not be—as the constitutivist would have it—a conceptual claim concerning the constitutive link between the act-type of assertion and its norm. Further, it is still not clear what would give this observation normative authority over the few speakers that do not follow the alleged norm of assertion. The fact that the Totos'

¹³ In experimental work, for instance, John Turri (2016) has shown that a majority of participants considers knowledge to be the norm of assertion. Crucially, his data also shows that there is a good number of people who disagrees.

wedding rituals in Bhutan are performed by the world's smallest tribe by no means makes them the wrong kind of rituals.¹⁴

But, more importantly, even if it were true that the actual practice of assertion is governed by a single communal norm across all actual linguistic communities, this would still fail to substantiate the constitutivist claim. Presumably, assertion is fundamental to our communicative practices, across languages, traditions, and communities. Under a constitutivist picture, the norm of assertion individuates the unified and generalized act-type of assertion, no matter how it is performed. Assertions can be performed by Spanish-, English-, and Klingon-speaking communities—and even by communities that we are yet to imagine. Hence, if the norm of assertion constitutively individuates the act of assertion, it should be able to adjudicate what counts as proper and improper assertions across communities, even possible ones. This means that a constitutive link between assertion and its norm would individuate assertion across worlds.¹⁵ After all, **Constitutivism** holds that the norm of assertion *essentially* individuates the unique norm of assertion.

Yet, the contents of communal norms are determined by the *contingent* stories and attitudes of the communities that generate and enact them. The development of a given game, convention, social norm, or e-function is a contingent matter, explained by the specific interactions within a community and in its cultural and social environment. Even if the actual communal norm of assertion was socially selected to provide a single epistemic good—say, justified belief—it is still possible for norms with different contents to be generated and selected at other possible times, cultures, and worlds.¹⁶ Hence, if the actual norm of assertion is a communal norm, it would still lack the normative authority to adjudicate whether it or other *possible* norms constitute the unique act of assertion across worlds.

At this point it may be tempting to just drop **Constitutivism** altogether. Perhaps there is no single norm of assertion that essentially individuates the unique act of assertion. Perhaps more than one norm can (or even actually does) govern our assertoric practices. The rest of

¹⁴ Moreover, if all there is to the norm of assertion is that humans contingently came to develop the game of subjecting our sayings to the ϵ -norm and not the ϵ^* -norm, it seems more plausible, as Cappelen (2011) argues, that there is a plethora of assertoric practices, none of which has the normative authority of being called *the* act of assertion. After all, as anthropologists and sociologists would argue, communal norms are notoriously diverse and variable across communities.

¹⁵ Note that this argument differs from the objection to conventionalism about assertion that García-Carpintero (2019) aims to defuse. That objection holds that assertion cannot be conventional because conventions are not normative across worlds. My point is that, even if a given convention is normative across worlds, nothing about that single convention would have the normative authority to adjudicate between itself and alternative conventions that could govern the practice of assertion across *other* worlds.

¹⁶ Simion & Kelp (2020) acknowledge that e-functionalism cannot fulfill the task of essentially individuating the act of assertion. Hence, they call their view 'light' constitutivism. It is not clear to me, however, how abandoning the essentialist claim of constitutivism still counts as constitutivism. So, as I've noted before, I will postpone the task of arguing why cooperativism is better than an e-functionalism that does not accept constitutivism as I have construed it here.

my argument shows that things are not so dire for the constitutivist. Replacing communalism for an instrumentalist picture based on the practical rationality of cooperative action yields a much better claim for the constitutivist about assertion.

To anticipate my argument, consider an analogy. Imagine that you and I are standing in the parking lot, and you ask me to give you the car keys. Say I accept the request but throw the way keys off, by about ten feet. There is something wrong with my act. I failed at throwing you the keys. Moreover, it seems that I have failed *you*; you can reproach me for my defective throw. This suggests that there is some norm that governs my action.

Nevertheless, if we want to understand my action and its normative profile, it would be theoretically idle to posit a communal norm that socially constitutes the practice of throwing keys. There is no preestablished, communally constituted practice of throwing keys that we have to learn in order to throw keys—or throw things, for that matter—in a way that is criticizable by others. Rather, the normativity of my act seems to be explained by the fact that I failed to do what I committed to do *with* you. Of course, if throwing keys becomes an important part of our communal living, we may develop communal norms that guide, constrain, and motivate key-throwing. Further, we may come up with informative stories about how we developed the means to throw keys at each other. These norms and stories, however, will not tell us what my act of throwing the keys across the parking lot amounts to. The rest of this paper argues that assertion is no different.

3. Assertion as Cooperative Action

Let us start with a Gricean/Lewisian insight. Communication often occurs when speakers embed their actions in a particular kind of rational interaction. Specifically, paradigmatic declarative acts—what we will call assertion—stem from a distinctive kind of coordination game in which agents acknowledge they have a common interest that is best served by interlocking their actions in cooperation.^{17, 18}

Consider an example. We are mutually and responsively acting (i.e., cooperating) towards planting a tree in my garden and it is clear to both of us that we need a shovel to dig a hole. I utter “behind the shed”. To make sense of my action as guided towards achieving our mutual interests, it must be explained as communicating that the shovel is behind the shed.

¹⁷ This claim is widely accepted across disciplines, including, philosophy, linguistics, psychology, and computer science. See, for instance, Asher & Lascarides, 2013; Cohen, 1982; Cohen & Perrault, 1979; Grice, 1991; Lewis, 1975, 1980; Litman & Allen, 2003; Geurts 2018; Roberts, 2012, 2018; Simons, 2018; Stalnaker, 1984, 2014.

¹⁸ Note the similar starting point of Greco (2022) account of testimony and cooperation. While there is much we agree on, Greco does not take the norm of assertion to be constitutive of assertion as I have construed it, as he allows for various norms of assertion with different contents. Hence, our accounts of the basic normativity of assertion differ—although mostly in complementing ways. However, since I am focused with developing a constitutive account of the norm of assertion, I postpone discussion of the divergence of our views. See also §5 for a distinctive derivation of what Greco calls the different norms of assertion.

As a speaker, I tailored my utterance so that you would interpret it as actively advancing our goals by communicating that the shovel is behind the shed. In turn, your interpretation of my utterance as communicating a given content is warranted by the presumption that it plays an instrumental role in advancing our mutual interests.

Altogether, what makes paradigmatic assertions a distinctive kind of signal is that the reasons why a speaker produces it and the reasons why her audience interprets it hinge on the fact that they share a common interest and they are mutually and responsively acting to achieve it. Cooperation serves as the interpretive fulcrum between a speaker's reasons to issue a communicative act and the audience's reasons to interpret an act as communicative.

To assert, I suggest, is best captured as the act of presenting oneself as cooperatively communicating information to advance the purposes of the conversation. More schematically:

Assertion: a speaker's utterance u amounts to asserting p whenever u must be interpreted as aimed to inform that p to rationalize u as a cooperative move in a conversation (i.e, as advancing the shared goals of the conversation).

To make sense of my utterance "it's behind the shed" as meant to advance our shared goals, it must be interpreted as communicating that the shovel is behind the shed.¹⁹ Assertions are distinguished from sarcasm, pretense, mere recitation, and even sleep-talking by the fact that they are utterances whose communicative function is given by their instrumental role of providing information to advance the shared purposes of the conversation. Presumably, this involves having some kind of epistemic authority over what one asserts. However, determining what this authority involves amounts to answering the normative question of assertion. And that will not come until §6. For now, I remain neutral on this issue.

Note that, although Grice-inspired, the cooperative view does not give the speaker's intention-in-action a fundamental role in determining the content of an assertion. What content is asserted by a token assertion is determined by the relation of an issued utterance with the mutually acknowledged shared goals that guide a conversation in which this utterance is embedded. This means that assertions can diverge from the specific intentions with which the speaker issued her utterance. In the case above, if you and I are looking for the shovel, and I say, "behind the shed," I assert that the shovel is behind the shed—even if for some reason my intention-in-action was to communicate that the book I was reading is behind the shed.

Let me sharpen this cooperative view by considering two worries about its extension. First worry: one may think that assertions can occur in the absence of cooperation—as in cases

¹⁹ Whether an action must be interpreted as intentionally tailored as a cooperative move will depend on considerations concerning the coherence, stability, and simplicity of an intentional action, given an acknowledged goal, see Bratman (2014), Bratman et al. (1988), Harman, (1990), Cohen & Levesque (1995).

of deception or cases of contentious communication in which interlocutors' interests are not fully aligned.

Note that the sense of cooperation at stake here is capacious. Two agents cooperate when their actions are guided by their mutual acknowledgment that they are interdependently acting towards the same outcome. That is, all that is required to cooperate is that agents mutually acknowledge that they are in a coordination game where the payoffs are aligned towards the same goals.²⁰ These goals, of course, can vary widely. Assertion may be enabled by non-communicative goals, like fixing a garden. Yet the goal itself may also be to simply share relevant information, as when you begin a conversation by asking me where Leonard Cohen is from.²¹

This means that cooperation can ensue in less-than-ideal cases. You may be motivated to fix the garden with me because of our friendship. Yet, you may also be motivated by the threat of punishment if you do not help me. Although agents are often motivated to not cooperate by staying silent or lying, a whole set of social (often unjust) devices and sanctions are often put in place to counter these motivations and promote the acceptance of shared goals that enable cooperative communication.²² Assertions can still occur in such contexts. By the same token, communicative cooperation can also be embedded in mostly antagonistic relations. Two war generals may be committed to killing each other in battle. Nonetheless, they may be able to cooperate in communication, as they still share the goal of finding a proper battlefield.

Further, acknowledged shared goals need not be personally held individual goals. An interlocutor may present herself as committed to the goal that guides the conversation but also fail to be genuinely committed to it. Deception is a clear example. My utterance counts as asserting that the shovel is behind the shed against the mutually acknowledged goal of planting the tree, even if in fact I know the shovel is not there. For my deception to be effective, however, I must give you good indication that I am indeed committed to our shared goal. Similarly, assertions can occur when agents are uncertain whether their counterparts are genuinely cooperative. You may have doubts about whether I am genuinely committed to

²⁰ This allows for some classic cases of isolated utterances discussed in the literature—like diary entries or soliloquies (Lackey 2007; Owens 2006). Since the relation between interlocutors can be one of time-extended identity, as long as a speaker can be seen to be cooperating with her future or past self, her utterances may count as assertions. If I know that Raul's keeps his diary to log relevant information for his future tasks, his entry "the drain is clogged" will count as an assertion. Otherwise, isolated utterances are derivative uses of declarative force.

²¹ An unprompted utterance can also have a cooperative role by implicitly accommodating new goals that are to be shared by interlocutors. Just as my asking you for directions can count as inviting you to share the shared goal of sharing information, my uttering "it's on the other block" can be accommodate the invitation to share information about the restaurant you are looking for. See Roberts (2018), Cohen & Levesque (1988), Asher & Lascarides (2005), for formal accounts of accommodation.

²² I come back to this dynamic in the discussion of the role of social norms of assertion in §5.

fixing the garden. Nonetheless, given that I have presented myself as committed to doing so, my utterance is still contextualized as an assertion (yet, perhaps one you should not believe).

Second worry: one may think that some cooperative communication will not count as assertion—as in some insinuations in which speakers may cooperatively communicate without asserting. Note, however, that while conversations require some degree of cooperation (after all, interlocutors with no shared interests are not motivated to even begin a conversation), not everything we do in conversation amounts to conveying a single cooperative interpretation. Most insinuations, for instance, happen in interactions where cooperative conditions are uncertain, and speakers aim to maintain some deniability for what they communicate (cf. Pinker 2007; Asher & Lascarides 2016; Camp 2019).

When asked by our boss, I may insinuate that my friend Gale bribed the police by uttering “he solved it like a businessman”. Given that my boss has reason to think I am not particularly motivated to sell out my friend, he knows that I am not committed to a shared goal that would render my utterance as cooperatively communicating that Gale bribed the police. Hence, my utterance is cooperatively compatible with this content *and* with other innocuous interpretations, like that he negotiated with noncriminal reasons. So, while my utterance may play some cooperative role in the conversation, it does not play the role of communicating *a* cooperatively required content—i.e., it does not assert that Gale bribed the police.²³ Insinuation requires cooperative ambiguity. Assertion precludes it.²⁴

My aim here has been to characterize the paradigmatic, normatively significant act of assertion. We have seen that it captures an intuitive range of assertions and non-assertions. Nevertheless, the paradigmatic approach here allows for borderline cases in which speakers may be taken to assert in the absence of cooperation—even in the permissive sense delineated above. Importantly, however, these borderline cases can be understood against the characterization above and, given their unclear cooperative role, will be normatively deficient in some way; that is, failed borderline assertions will not clearly breach the normative profile of assertion.²⁵ I turn now to characterizing this profile.

²³ Of course, insinuations may occur in less contentious interactions. Perhaps my boss and I are just gossiping about our officemates. Even in those cases, however, if my utterance is to play the deniability-preserving role of insinuation it must be vague enough to allow for equally good interpretations given the acknowledged goals of the conversation.

²⁴ See Vesga (forthcoming) for an account of testimony on these lines.

²⁵ Consider an example. You and I are watching a movie I have seen before. You hate spoilers. Yet, I utter the “that guy is a ghost.” One may think my utterance still counts as an assertion, even if quite literally, your motivations are against having that information. Under the cooperative view, a) this counts as a borderline assertion because I knew you would contextualize my act as an assertion. I acted *as if* we were cooperating and thus forcefully accommodated your interpretation of my act as an assertion (see n. 19). And, crucially, b) my declarative act would have a defective normative profile. There is something normatively suspect about imposing unwanted information. If the character in question was not a ghost, it is not so clear that I have failed you in the same way I would have if I offered a defective paradigmatic assertion—e.g., if there is not a restaurant in fifth and Cascadilla in our opening example above.

4. The Instrumental Norm of Assertion

A speaker asserts whenever she presents herself as cooperatively communicating that p in order to advance the shared purposes of the conversation. This means that assertion is characterized by the epistemically significant aim of providing information as an instrumental move within a cooperative engagement.

This view of assertion follows a particular sense of game—one closer to game theory than to chess. Two agents engage in a coordination game when they mutually acknowledge the relevance of each other's actions to achieve their goals. In this sense, games do not depend on previously determined rules that constitute them. All it takes is two or more agents who strategize their actions, given that the other is doing the same. In turn, assertions stem from a distinctive kind of coordination game in which agents mutually acknowledge that their common interest is best served by interlocking their actions in cooperation.

A direct consequence of characterizing assertion as enabled by the practical rationality of cooperation is that failing to provide the information one asserts is an instrumental failure. A speaker is able to assert p because she presented her utterance as offering information as means to achieve a shared goal. Therefore, failing to indeed provide the information one presents oneself as offering amounts to an instrumental failure to act toward the goal that structure the conversation—just as passing you baking soda instead of salt is an instrumental failure to achieve our joint project of making pasta. If it turns out that the shovel is not behind the shed, I have failed to act as my cooperative commitment to fix the garden required. I have issued a defective assertion with respect to the aim I presented it as having.

This means that under the cooperative view, mutually acknowledged cooperative commitments to act together towards the same goal set the standard against which assertions count as appropriate or inappropriate. Given that an asserting speaker presents an utterance as aimed at cooperatively offering information, said utterance is subject to the instrumental normativity of this aim. This, I submit, is the normative core of assertion and allows for parsing the following instrumental norm.

Cooperative Norm of Assertion (C-Norm): If a speaker is to cooperatively offer information as means to achieve a shared goal, they must actually provide said information.

Note that at the very least this amounts to an instrumental requirement of true assertions. If one presents an utterance as having the aim of providing the information that p and p is false, one has failed to achieve this aim. We will see in §6 that the normative requirements of this instrumental norm are arguably much more robust. However, since my focus here is on the *metanormative* profile of this instrumental requirement, I postpone this issue.

Under this construal, the aim of cooperatively offering information is constitutive of assertion. To even be able to assert, a speaker must present herself as committed to the aim of cooperating towards the shared goal that structures the conversation and contextualizes her utterance as an assertion. And, in turn, this commitment sets the standard of success for her assertion. Put otherwise, assertions are subject to the instrumental pressures entailed by the cooperative commitments that enabled them in the first place. There is no assertion without presenting an utterance aimed to cooperatively offer information. Hence, there is no assertion that is not subject to the instrumental norm this aim entails.²⁶

The cooperative view has a straightforward answer to what Ishani Maitra (2011) calls the *hard question* about assertion (p. 279): what kind of mistake does a speaker make when she issues a defective assertion? Under the cooperative view, defective assertions are instrumental failures to achieve the shared goals of the conversations in which they are embedded.

The cooperative norm also makes sense of how assertions can be defective and still be assertions. A defective assertion amounts to an utterance that was presented as offering information as a means to achieve a shared goal, but fails to actually provide said information. In that case, a speaker issues an assertion but fails to comply with her acknowledged commitment to providing information that would advance the shared goal that structures the conversation. She fails to fulfill the cooperative aim that contextualized her utterance as an assertion in the first place.

In turn, a speaker who is clearly not willing to cooperate at all is not in the position to assert. Such an agent is either not interested in even beginning to have a conversation or will simply not be interpreted as issuing an assertion.²⁷ This is why the Moorean-style move of saying “p, but I don’t $\epsilon(p)$ ” is often heard as non-sensical or as shifting the kind of act being performed. Under the cooperative reading, it amounts to presenting the information that p as means to a shared goal but disavowing that the audience can count on that information to achieve said goal. Asserting, “Tomorrow they serve meatballs, but I have no idea if they do” is instrumentally absurd once you present yourself as a cooperative interlocutor interested in

²⁶ The assumption that there are mean-ends requirements of instrumental or procedural rationality is unproblematic. Why this is the case and the logical structure of these requirements is not so. All the argument here needs is the narrow scope reading that, given commitment to a goal, one is committed to achieve the means necessary to achieve it (at least in a minimally procedural sense.) For a thorough review on this subject, see Kolodny & Brunero (2023).

²⁷ What about bald face lies (Sorensen 2006)? If the cooperative view of assertion is right, I am inclined to say bald-face lies are not proper lies. A lying speaker asserts p with the clear intention of making her audience falsely believe p. To do this, the speaker must present herself (albeit in bad faith) as a cooperative speaker. The failure of lying is therefore a betrayal of a cooperative commitment. If I issue the bald-face lie “that house is green” when it is clear to both of us that it is blue, it is not clear that I am subject to the same kind of criticisms to which the lying speaker is. I am not betraying your trust because I am clearly not interested in giving you useful information. For a constitutivist who thinks that the norm governing assertion is essential for its ontology, this should be a crucial datum.

offering information as means to achieve a shared goal. At most, this counts as guessing or speculating—not asserting—that tomorrow they will serve meatballs.

Note too that, as opposed to failures in perception, memory, or belief, defective assertions have a distinctive second-personal character. When I assert that there is a restaurant on fifth and Cascadilla when there is not, my failure essentially involves you in ways that entitle you to reproach me. Once you and I commit to do φ together, acting in a way that detracts us from φ is not only a failure to do what I committed to do, but a failure to do what I committed to do *with* you.²⁸ Given this mutual commitment to a goal, you are entitled to act in ways that preserve this commitment, which include reproaching me or holding me accountable if I diverge from our commitments.²⁹

Finally, the cooperative view is not subject to an objection commonly raised against epistemic constitutivism in general. As the argument goes, constitutivism may be able to explain why a phenomenon φ essentially involves a norm n but not why anyone should follow the norm itself (Côté-Bouchard 2019; Enoch 2007; Wedgwood 2009). We may imagine a constitutive account of Baroque etiquette that involves norms of gender asymmetries in serving food. Nevertheless, this account has not offered an account of why anyone should follow these norms altogether. Similarly, while constitutivism may account for a norm that prescribes following certain epistemic norms, it does not explain why such norms should be followed.

Nevertheless, I have not offered a view of why agents should assert in general. The present view holds that once a speaker asserts—once she has presented herself as cooperatively offering information—she is subject to epistemically significant instrumental requirements. The instrumental norm of assertion governs only those who present themselves as cooperative. Hence, the constitutivist picture here is, as Luca Ferrero (2019) would put it, a normatively modest one. Once agents have reason to commit to a shared goal and provide information as means to achieving it, they have reason to comply with the standards set by this aim. We will see in the next section how the modesty of the instrumental norm of assertion helps to explain the non-constitutive role of communal norms of assertion.

5. Constitutive and Non-Constitutive Aptness

Before turning to the normative task of specifying what the norm of assertion epistemically requires from speakers, let me note two further explanatory merits of the cooperative view.

²⁸ The normativity of assertion is often compared to that of promising (McMyler 2011; Fricker 2004; Moran 2018). This parallel would be maintained if promising is a commitment in joint deliberation, as Brendan de Kenessey (2017) has it.

²⁹ Of course, if I committed to cooperate in bad faith and I am not personally committed to actually cooperate, your reproach does not have much purchase on me. Yet, from the perspective of the cooperative commitment that enabled me to assert, I have entitled you to reproach me if I fail to give you the information I presented myself as offering. See Vesga (forthcoming) for an argument in this line when it comes to testimony.

First, I address how it avoids the adjudication problem presented in §2. Second, I address how it provides a unified picture of the non-constitutive role other conventional, social, game-like, and functional norms have with respect to the act of assertion.

For starters, note that the cooperative view is not subject to objections raised in Pagin (2016). There is no need to posit a communal norm that makes up the act of assertion and is generally, but also implicitly, accepted across every linguistic community. Nor it is necessary to posit a cognitively complex attitude of knowing (without necessarily being able to externalize) the communal norm of assertion. Under the cooperativist view, knowing the norm of assertion amounts to knowing how to assert. As with our key-throwing example, once we see that assertion can be individuated by its aim within a cooperative scheme, it would be theoretically idle to posit a further meta-grammar of socially constructed rules that enable one of our most basic speech acts.

More to the point, just as ‘driving’ is a rigid designator of the act of driving, no matter the means or vehicles with which it is performed, ‘assertion’ rigidly designates the same act across worlds, no matter the specific means by which it is performed or the specific histories of how agents came to be able to perform it. Of course, we developed cars, buses, and motorcycles to better drive. Similarly, we developed Spanish, Thai, and Esperanto to better assert. But it is vital to distinguish an account of what it is to drive from an account of the different means of driving. Similarly, it is important to distinguish an account of what it is to assert from an account of the means of assertion. And it is the former task the constitutivist is interested in.

Under the cooperative view, as long as an utterance is to be interpreted as cooperatively offering information in a coordination game, said utterance amounts to an assertion, subject to the instrumental norm that this aim entails—no matter the means by which this assertion is performed. Hence, the nature of assertion is not contingent on the histories or practices of specific human communities. Even speakers whose cognitive abilities, social practices, and linguistic repertoires owe to very different developmental and evolutionary stories will assert as long as they use these means to cooperatively communicate information.

Note that this is only the metanormative story. I have not yet said much about what the act of assertion epistemically requires of the speaker. In the next section, I will argue that the link between success and cooperative action suggests that assertion requires speakers to have a distinctive epistemic standing with respect to the content of their assertions that is at least close to knowledge. You may disagree. Perhaps assertion requires only justification or truth. In any case, under the cooperative picture, our disagreement is substantive. Solving it will require specifying the constitutive aim that individuates assertion as an act-type across worlds, which will in turn adjudicate our disagreement.

Before I throw my hat in the normative ring, however, let me turn to a final metanormative advantage of the cooperative view. Not only does it yield an apt view of the constitutive status of the norm of assertion; it also yields an apt view of the crucial-yet-not-constitutive role other norms have in guiding, motivating, and constraining the act of assertion.

5.1 Conventions. Under the cooperative view, assertions are partial solutions to coordination problems. The motivations of interlocutors align towards a shared goal, which enables them to issue and interpret asserted contents as means to achieve said goal. Crucially, when agents continuously face a coordination problem, they often establish conventions to solve it across times and contexts. If you and I want to meet for coffee every Tuesday, developing the convention of meeting at Give'Ya Coffee at ten in the morning is an effective way to solve our coordination problem. Languages are conventions in this sense (Lewis 1980; Marmor 2009; Geurts 2018).³⁰ Faced with the recurrent problem of sharing information to achieve shared goals, we develop arbitrary linguistic patterns to share information more effectively. Crucially, then, specific conventions can serve as arbitrary and ever-developing means that guide (but not constitute) the act of assertion, such that it can be communally performed more effectively among members of a given linguistic community.

5.2 Social Norms. As I mentioned in the last section, the cooperative view yields no requirement to cooperate. Rather, it characterizes the communicative requirements for agents who already hold a mutually acknowledged commitment to cooperate with each other. However, agents are motivated to avoid cooperation altogether or even to present themselves in bad faith as committed to a shared goal to which they are not personally committed. In that case, communicating agents face a mixed motive game where their goals are not fully aligned in a way that enables actual cooperation. A parallel scenario occurs when guests of a potluck are interested in getting to the food before anybody else. From the individual perspective of each agent, cooperation is not properly motivated. Yet, faced with mixed motive games, communities develop social norms—sets of informal sanctions and social expectations—that motivate agents into cooperation (Bicchieri 2006, Brennan et al. 2016; Elster 1995). These communal sanctions and expectations motivate agents against selfish, uncooperative actions. Faced with the cost of social reprimand, for instance, guests are motivated to cooperate into forming a queue to get food. Similarly, linguistic communities develop social norms sanctioning liars or irresponsible speakers that motivate interlocutors to effectively cooperate in contexts in which they would not do so otherwise (Graham 2020, Starr 2018). Social norms

³⁰ See Millikan (1987) for a different account of convention. *Mutatis mutandis*, the cooperative view can also explain the non-constitutive role of assertion conventions under this parsing.

of assertion turn mixed-motive contexts into coordination games that enable honest assertions. However, speakers whose motivations naturally align are perfectly able to assert in the absence of these norms. Hence, these social norms of assertion motivate but do not constitute the act of assertion.

5.3 Game-Like Norms. We have seen that the cooperative view is parsimonious. It does not require the cognitively odd posit of an implicitly accepted but never acknowledged communal norm in people's minds. It only requires instrumental reasoning and abilities conferred by our cooperative cognition (Tomasello, 2016; 2012). However, communities often develop further rules and procedures that specify the conditions under which assertions (false or true) may have clear distinct institutional consequences. This is the case in court testimony, for instance. Different legislatures develop norms that specify the conditions under which some acts count as court-given testimony and the consequences of offering it. Similarly, communities may develop explicit rules that, in certain circumstances, bestow assertions with particular socially significant consequences. These rules constrain the practice of assertion in ways that generate derivative assertoric practices. Nonetheless, no such rule is necessary to perform the act of assertion. Hence, they do not constitute the act of assertion proper.

5.4 Functions. As with biological traits, etiological functions dissipate the mystery of why certain social practices are subject to normative standards when they were not intentionally designed to have them. Some communities, for instance, may have non-intentionally developed functional norms of marriage because this practice ensured property arrangements and thus their thriving. Yet, we have seen, etiological functions are contingent to the interaction of the target phenomena and their environment. In some other communities, marriage norms may have developed because it eases childbearing. Hence, when it comes to assertion, etiological functions can, at most, explain why certain communities' assertoric practices have certain features and not others. They will not explain what assertion is or what it demands of speakers across actual and possible worlds. Nonetheless, as Neander (1991) notes, there are certain functional features whose normativity is no mystery: the functionality of intentionally made artifacts. A rock has the *artifactual* function of hammering if I use it with the purpose of hammering. No ancestor needed. My selecting the rock for this purpose imbues it with the function of hammering. Likewise, token utterances are assertions because they are selected by a speaker to fulfill the constitutive aim of assertion. Hence, the cooperative view can be rendered as an *artifactual*—as opposed to etiological—functionalism, such that the functional norm that uniquely individuates the act of assertion is derived from the constitutive aim of the act of assertion.

Summing up, the norm of assertion is determined by the aim that characterizes the act of assertion. If a speaker is to cooperatively offer information, they better have it. Given this normative profile, assertion can also be governed by further norms that guide, motivate, and constrain the conditions under which interlocutors cooperate in communication. These norms are highly important in explaining and ensuring large-scale coordination and communal cohesion. We have seen, however, that they are contingent to the histories and selection processes of specific communities and lack the normative authority to adjudicate what counts as the act of assertion proper.

6. Deriving The Normative Story

We have a metanormative story of assertion. Assertion is individuated by the aim of cooperatively offering true information as means to achieve a shared goal. And it is in virtue of the instrumental constraints derived from this aim that assertion is subject to its norm. In this last section, I explore how this metanormative story suggests a novel path to a normative story about the content of the norm of assertion. That is, I explore here whether the answer to *why* asserting speakers are subject to the norm of assertion can yield an answer to *what* this norm asks of them.

My main purpose here has been to motivate a distinctive metanormative account of the norm of assertion. Nevertheless, I take it as a crucial advantage of the cooperative view that it offers a novel strategy to address the normative task of deriving the content of this norm. So, although a full defense of this normative account is beyond the scope of this paper, I chart a plausible argument for thinking that the norm of assertion requires at least something close to knowledge. Whether it is knowledge, however, will come down to the reader's epistemological allegiances.

The key move is to note that intentional action has a distinctive practical profile, which is shared by the epistemically significant act of assertion as a special case. To begin, note that there is a strong link between intentional cooperative action and success. Recall our key-throwing example. You ask me to give you the keys and I accept. One understanding of being cooperative is to merely *try* to be cooperative. All it takes to be cooperative in this sense is to make an honest effort and have the intention of throwing you the keys—even if I throw them way off. But in a stronger sense, I act cooperatively when I act in a way that actually brings about our shared goals. I am cooperative in this sense only if you actually catch the keys.

So, we may distinguish between *acting cooperatively towards* φ and a *cooperatively* φ -ing. The former allows for honest efforts and unsuccessful attempts. The latter is only the case when one successfully (and thus actually) performs φ . Hence, we may posit the following plausible principle:

Actuality: An agent cooperatively φ -s only if they actually φ .

If I throw the keys way off so that you cannot catch them, I have not cooperatively given you the keys because I simply did not give you the keys. Crucially, when one acts cooperatively towards φ one aims to cooperatively φ . If I act cooperatively towards baking a cake with you, I aim to cooperatively bake a cake with you—even if I ultimately fail to do so.

Note, however, that cooperatively φ -ing requires more than just actually φ -ing. There must be some belief constraint on intentional cooperative action. To intentionally throw the keys at you at time t , at the very least, I must not believe that it is impossible for me to throw the keys to you (Bratman 1995, 2001). Yet, it seems too strong to say that to cooperatively φ one must believe that one will φ *simpliciter* (Goldman 1970; Harman 1976). So, a plausible and general constraint on cooperatively φ -ing is that one must believe that one is sufficiently likely to φ .

Belief: An agent cooperatively φ -s at time t_2 only if she believed at t_1 that they are sufficiently likely to φ .³¹

How to cash out this likelihood belief is a complicated matter. But for our purposes, what matters is that if one cooperatively (and thus intentionally) φ -s, one must fulfill this plausible doxastic requirement on intentional action.³²

Cooperatively φ -ing needs more than **Actuality** and **Belief**. Say that I throw the keys way off. Yet, unbeknownst to me, they hit the one bird that was flying around, causing the keys to fall on your hand. In a sense, my action is causally connected to what made our shared goal true. Yet, there is not a proper link between my agency and my actually achieving our goal. We may thus posit a third plausible claim:

Agency: An agent cooperatively φ 's only if there is a proper connection between her agency and her φ -ing.

Of course, more needs to be said about what this proper link precisely amounts to. And this is no trivial matter. Yet, for our purposes, we can say that cooperatively φ -ing precludes deviant causal chains between my intention to φ , my bodily movements, and my actually φ -ing.

Crucially, this proper link between my agency and the achievement of our goals must be modally robust. Say that part of my plan of getting you the keys indeed involves throwing them randomly to the sky and hoping that a bird passes and hits the keys so that they land on your hands. This is not acting cooperatively—even in the astronomically lucky case in which

³¹ See Bratman (1987), Pavese (2021).

³² I am partial to cashing this probabilistic belief as a sufficiently strong credence that one will φ . See below, and Moss (2019) and Greco (2015) for considerations in favor of this being enough for having probabilistic knowledge.

it actually happens. The problem with this scenario is that I acted irresponsibly to achieve a goal I committed to achieve with you. I did not have adequate control over the outcomes of my actions. If it ensues, I was just (extremely) lucky. Hence, cooperatively ϕ -ing requires that my actions lead to the intended outcome in sufficiently similar cases—such that it could not have easily been the case that I did not ϕ .³³

Although I will return to the question of how modally robust cooperative success must be, these considerations suggest that cooperatively ϕ -ing requires acting in a way that would lead to sufficiently robust success. This allows for a fourth plausible claim:

Robustness: An agent cooperatively ϕ -s only if she would have ϕ -ed in a sufficient range of similar cases.

Just as one cannot intentionally win a fair lottery³⁴, I cannot cooperatively give you the keys by throwing them way off.

Summing up, once an agent commits to a shared goal that requires cooperatively ϕ -ing, she commits to ϕ -ing successfully, which involves i) her actually ϕ -ing, ii) a doxastic attitude about one's ϕ -ing; iii) a proper link between her agency and her ϕ -ing, and iv) her actually ϕ -ing being modally robust. Of course, we often fail to act cooperatively in this sense. And we often have good excuses for our failures. What matters, however, is that when we commit to cooperate, the standard against which my actions count as failures is given by the aim of cooperatively ϕ -ing, which involves these features.

Now let us return to assertion. Recall that under the cooperative picture, a speaker asserts when she presents herself as cooperatively offering information. You and I are looking for a restaurant and I utter “There’s a restaurant on Fifth and Cascadilla.” I have presented myself as cooperatively offering the information that there is a restaurant on Fifth and Cascadilla as a means to achieve our shared goal. Whether my assertion is defective or not is determined by whether it fulfills the aim I presented it as having. In turn, given that assertion is an epistemically significant act, the practical profile of intentional cooperative action entails an epistemic profile of assertion.

Given **Actuality**, if I am indeed cooperatively offering the information that there is a restaurant at Fifth and Cascadilla, it must be the case that there is indeed a restaurant at Fifth and Cascadilla. If there is no restaurant, I have not cooperatively offered information; and I have failed to do what I presented myself as doing.³⁵ Hence, to be successfully cooperative, my

³³ Compare to Mele (1995) and Mele & Moser (1994).

³⁴ Gibbons 2001; Pavese 2021; but see Cath (2015). Compare too with Shepherd’s (2021) modal of agential control.

³⁵ One may object that sometimes providing false information may be cooperative. Say that Ole wants to steal the dossier that is in Jody’s office and Tan offers to help. Tan tells Ole that the dossier is on the first drawer on the right. However, Tan knows that the office Ole is going to is not Jody’s, but Fran’s. In that sense even if the in question dossier is not in Jody’s office, it

assertion must be true.³⁶ In turn, **Belief** suggests that if I cooperatively offered the information that p at time t, I must have believed that I was sufficiently likely to actually do so. This means that I must at the very least have believed that there is a sufficient likelihood of there being a restaurant at fifth and Cascadilla.

Further, given **Agency**, cooperatively providing information also requires a proper connection between the fact that there is a restaurant on fifth and Cascadilla and my *epistemic* agency. I must have good reasons to believe what I say. If my belief was mere wishful thinking or if I was just guessing, I was not actually cooperative. My epistemic agency (be it understood in terms of my access to evidence or my sensitivity to epistemic reasons) must be properly linked to the fact that I provided information that led to our goal. I must have a good epistemic standing—i.e., justification—to believe what I said. And, moreover, given **Robustness**, my cooperatively offering information requires this link to be modally robust. If I was being actually cooperative, it cannot be a mere fluke that there is a restaurant on Fifth street. There must be a restaurant in sufficiently similar possible worlds in which I asserted this content. That is, my assertion must be safe.

Hence, when it comes to the epistemically relevant act of cooperatively offering information, the standards of intentional cooperative action require the speaker to have a particular epistemic relation to said information. Specifically, the standard of assertion requires truth, belief, justification, and safety.³⁷ That this epistemic profile is at least entailed by knowledge is unproblematic. Whether it amounts to knowledge, however, will depend on

may be said that Tan cooperatively offered information that advanced their shared goals. Nonetheless, Tan was *not* cooperative with respect to the goal he presented himself as sharing with Ole (stealing Jody's dossier). The information he gave is not a means to steal Jody's dossier, simply because that is just not Jody's dossier. Even if under some description of his action, Tan ended up helping Ole achieve his own goals, Tan was not mutually and responsively acting with Ole towards the same goal. Agents can be said to be genuinely cooperatively ϕ -ing only if they are doing so under the same description.

³⁶ This, and considerations similar to that in n.36, gives us reason to deny that speakers can properly assert what they believe to be false (*contra* Lackey 2010). Assertion amounts to the act of presenting p as true such that having said information is a means to achieve the goals that structure the conversation. If a speaker believes p is false, she cannot believe informing the truth of p will bring about the shared goals that structure the conversation; after all, she already believes p is false. Lackey's creationist teacher who disbelieves the assertions about evolutionary theory does not share with her students the goal of sharing true information—which is *ex hypothesi* what her students think she is doing with them. Hence, from the practical perspective of what she and her audience have publicly committed to do together, her assertion is not appropriate. She is not acting cooperatively towards the same goal as represented by her and her students. Additionally, this provides an error theory for the intuition that there is still something appropriate about her act, as from *her own* practical perspective she is still helping their students.

³⁷ This argument is similar to and inspired by other arguments concerning the connection between knowledge's epistemic profile and successful action—cf. Greco (2016), Williamson (2017), and Pavese (2018). Yet, note that my claim is limited to the act of assertion. When it comes to the epistemically relevant of cooperatively offering information, the practical requirements of successful action require a particular epistemic relationship between the speaker and the proposition she communicates. In other words, the fact that assertion demands a particular epistemic relationship to a proposition is already given by the fact that assertion is the act of cooperatively offering a proposition.

a substantive view of what knowledge is. Nonetheless, let me conclude this section by noting two final considerations.

First, and more conspicuously, **Belief** does not require full belief about what one asserts—only the probabilistic belief that it is sufficiently likely to be true. Initially, this may count against the idea that assertion requires knowledge. However, keep in mind that given the other conditions of successful action this belief must be ultimately true, justified, and safe. If, as has been strongly argued by Sarah Moss (2018), probabilistic belief can amount to knowledge, then it is open to the cooperativist view to hold that assertion requires probabilistic knowledge.

Second, a vital consideration here will be what counts as robust success in cooperative action—as it will determine the kind of safety required by the norm of assertion. One available option is to say that successful cooperative action requires successfully φ -ing in *all* close possible worlds. If so, assertion would demand what we can call *strong safety*: asserted contents must be true in all close possible worlds. This seems too demanding, though. Perhaps in this world you easily caught the keys. Yet there may be one close possible world in which a bird flies by and knocks them out of your reach. Still, it would seem I cooperatively and successfully gave you the keys in this world. It was not a fluke that you got them. Hence, it may seem that the requirement for robust success requires success in a sufficient range of close possible worlds. If so, assertion seems to require *weak safety*: asserted contents must be true in sufficiently close possible worlds.

Importantly, one may also think that the modal base by which a cooperative action may be deemed to be successful depends on the practical considerations of the case at stake. If we are going for casual dinner and I throw you the keys, accidents are not too costly. So, my acting successfully may allow for cases in which you actually catch the keys but in some close possible scenarios you do not. Yet, if our mutual friend is having an anaphylactic shock and we need to drive them to the hospital as soon as possible, I am plausibly required to act much more carefully—as accidents are much more costly. I must act such that there are less possible worlds in which you do not catch the keys. So, success in this case may require a more demanding standard of robustness. By extension, the practical stakes of a conversation may modulate the safety standard by which an assertion is successful.

Hopefully, it is clear how these considerations map onto the current landscape in epistemology. Some authors hold that only something like weak safety is required for knowledge.³⁸ Others hold that only something like strong safety can yield knowledge.³⁹ Yet

³⁸ Pritchard (2012), Sosa (2015), and John Greco (2016).

³⁹ Sosa (1999), Williamson (2000), Pritchard (2005), Manley (2007), and Daniel Greco (2016).

others argue for contextual safety standards for knowledge.⁴⁰ Again, this contention will be settled by further external arguments on the nature of knowledge itself.

Nonetheless, wherever these contentions land, the cooperative view of assertion brings the relationship between assertion and its required epistemic profile much closer together than communal constitutivism. The act of assertion does not demand its distinctive epistemic profile because a group of people came together and (implicitly or explicitly) generated and accepted a norm that demands it. Rather, the epistemic profile of assertion follows from a constitutive relation between the demands of successful cooperative action and the epistemic significance of the act of cooperatively offering information.

7. Conclusion

An adequate constitutivist account of assertion must address at least two tasks. On the one hand, it must address the normative question of assertion: it must offer an account of what the norm of assertion requires from speakers. On the other hand, it must address the metanormative question of assertion: it must offer an account of what kind of norm is the norm of assertion—or, in other words, *why* are asserting speakers subject to the norm of assertion.

I have dealt mainly with the latter metanormative question as an entry point to develop a full account of assertion. Constitutivism cannot hold that the norm of assertion is a communally generated, and thus contingent, norm. Instead, the constitutivist claim is best served by characterizing assertion as a cooperative action characterized by the constitutive aim of cooperatively offering information to achieve the shared goals that structure a given conversation. In turn, given the modesty of this constitutive account, the cooperative view offers a cohesive picture of non-constitutive norms of assertion and their crucial role in fostering the use of assertion in large-scale communal coordination. Moreover, I have suggested that if successful cooperation is the standard by which we judge cooperative action, there are good reasons to think that the epistemic profile of cooperatively offered information is at least entailed by knowledge.

References

Asher, N., & Lascarides, A. (2005). *Logics of Conversation*. Cambridge University Press.

Asher, N., & Lascarides, A. (2013). Strategic conversation. *Semantics and Pragmatics*, 6(0), Article 0.
<https://doi.org/10.3765/sp.6.2>

⁴⁰ Stanley (2005), DeRose (2008), Fantl & McGrath (2009), Stalnaker (2019), for instance, argue for contextual modulation of knowledge claims, although by way of significantly different paths.

- Benton, M. A. (2011). Two More for the Knowledge Account of Assertion. *Analysis*, 71(4), 684–687.
<https://doi.org/10.1093/analys/anr085>
- Bicchieri, C. (2006). *The Grammar of Society: The Nature and Dynamics of Social Norms*. Cambridge University Press.
- Bicchieri, C. (2016). *Norms in the Wild: How to Diagnose, Measure, and Change Social Norms*.
- Brandom, R. (1983). Asserting. *Noûs*, 17(4), 637–650.
- Bratman, M. (1987). Intention, Plans, and Practical Reason (Issue 2, pp. 198–199). Cambridge, MA: Harvard University Press.
- Bratman, M. E., Israel, D. J., & Pollack, M. E. (1988). Plans and resource-bounded practical reasoning. *Computational Intelligence*, 4(3), 349–355.
- Brennan, G., Eriksson, L., Goodin, R. E., & Southwood, N. (2016). *Explaining Norms*. Oxford University Press
- Camp, E. (2019). Insinuation, Common Ground, and the Conversational Record. In *New Work on Speech Acts*. Oxford University Press.
- Cappelen, H. (2011). Against Assertion. In J. Brown & H. Cappelen (Eds.), *Assertion: New Philosophical Essays*. Oxford University Press.
- Cath, Yuri. (2015). “Revisionary Intellectualism and Gettier.” *Philosophical Studies* 172(1): 7–27.
- Cohen, L. J. (1982). What Is Necessary for Testimonial Corroboration? *The British Journal for the Philosophy of Science*, 33(2), 161–164. JSTOR.
- Cohen, P., & Levesque, H. (1988). Rational Interaction as the Basis for Communication. *Pattern Recognition - PR*, 39.
- Côté-Bouchard, C. (2016). Can the aim of belief ground epistemic normativity? *Philosophical Studies*, 173(12), 3181–3198.
- DeRose, K. (2002). Assertion, Knowledge, and Context. *The Philosophical Review*, 111(2), 167–203.
<https://doi.org/10.2307/3182618>
- Douven, Igor. 2006. “Assertion, Knowledge, and Rational Credibility.” *The Philosophical Review* 115(4): 449–85.
- Enoch, D. (2006). Agency, Shmagency: Why Normativity Won’t Come From What is Constitutive of Action. *Philosophical Review*, 115(2), 169–198. <https://doi.org/10.1215/00318108-115-2-169>
- Fantl, J., & McGrath, M. (2002). Evidence, Pragmatics, and Justification. *Philosophical Review*, 111(1), 67–94.
<https://doi.org/10.1215/00318108-111-1-67>
- Fantl, J., & McGrath, M. (2009). Knowledge in an Uncertain World. In *Knowledge in an Uncertain World*. Oxford University Press.
- Ferrero, L. (2019). The Simple Constitutivist Move. *Philosophical Explorations*, 22(2), 146–162.
<https://doi.org/10.1080/13869795.2019.1599048>

- Fricker, E. (2004). Testimony: Knowing Through Being Told. In M. Sintonen, J. Wolenski, & I. Niiniluoto (Eds.), *Handbook of Epistemology* (pp. 109–130). Kluwer Academic Publishers.
- Fricker, E. (2006). Second-Hand Knowledge. *Philosophy and Phenomenological Research*, 73(3), 592–618.
<https://doi.org/10.1111/j.1933-1592.2006.tb00550.x>
- García-Carpintero, M. (2019). Conventions and Constitutive Norms. *Journal of Social Ontology*, 5(1), 35–52.
<https://doi.org/10.1515/jso-2019-0013>
- Geurts, B. (2018). Convention and common ground. *Mind & Language*, 33(2), 115–129.
<https://doi.org/10.1111/mila.12171>
- Geurts, B. (2019). Communication as commitment sharing: Speech acts, implicatures, common ground. *Theoretical Linguistics*, 45(1–2), 1–30. <https://doi.org/10.1515/tl-2019-0001>
- Gibbons, J. (2001). Knowledge in Action. *Philosophy and Phenomenological Research*, 62(3), 579–600.
<https://doi.org/10.1111/j.1933-1592.2001.tb00075.x>
- Graham, P. J. (2014). Functions, Warrant, History. In A. Fairweather & O. Flanagan (Eds.), *Naturalizing Epistemic Virtue* (pp. 15–35). Cambridge University Press.
- Graham, P. J. (2018). Social Knowledge and Social Norms. In *The Philosophy of Knowledge: A History. Volume IV: Knowledge in Contemporary Philosophy*. (pp. 111–138).
- Graham, P. J. (2020). The Function of Assertion and Social Norms. In S. Goldberg (Ed.), *The Oxford Handbook of Assertion* (pp. 725–748). Oxford University Press. <https://doi.org/10.1093/oxfordhb/9780190675233.013.30>
- Greco, D. (2015). “How I Learned to Stop Worrying and Love Probability 1.” *Philosophical Perspectives* vol. 29 (1): 179–201.
- Grice, H. P. (1967). Logic and Conversation. In P. Grice (Ed.), *Studies in the Way of Words* (pp. 41–58). Harvard University Press.
- Grice, H. P. (1968). Utterer’s Meaning, Sentence-Meaning, and Word-Meaning. *Foundations of Language*, 4(3), 225–242.
- Hawthorne, J. (2006). *Knowledge and Lotteries*. Oxford University Press.
- Henrich, J. (2016). *The Secret of Our Success: How culture is driving human evolution, domesticating our species and making us smarter*. Princeton University Press.
- Kelp, C., & Simion, M. (2021). *Sharing Knowledge: A Functionalist Account of Assertion*. Cambridge University Press.
- Kenessey, B. de. (2020). Promises as Proposals in Joint Practical Deliberation. *Nous*, 54(1), 204–232.
<https://doi.org/10.1111/nous.12269>
- Kukla, R., & Lance, M. (2009). ‘Yo!’ and ‘Lo!’: The Pragmatic Topography of the Space of Reasons. Harvard University Press.

- Kvanvig, J. L. (2009). Assertion, Knowledge, and Lotteries. In D. Pritchard & P. Greenough (Eds.), *Williamson on Knowledge* (pp. 140–160). Oxford: Oxford University Press.
- Lackey, J. (2007). Norms of Assertion. *Noûs*, 41(4), 594–626. <https://doi.org/10.1111/j.1468-0068.2007.00664.x>
- Lewis, D. (1975). Languages and Language. In K. Gunderson (Ed.), *Minnesota Studies in the Philosophy of Science* (pp. 3–35). University of Minnesota Press.
- Lewis, D. (1980). *Convention*. Blackwell.
- Litman, D. J., & Allen, J. F. (2003). Discourse Processing and Commonsense Plans. In P. R. Cohen, J. Morgan, & M. E. Pollack (Eds.), *Intentions in Communication*. The MIT Press. <https://doi.org/10.7551/mitpress/3839.003.0019>
- Maitra, I. (2011). Assertion, Norms, and Games. In J. Brown & H. Cappelen (Eds.), *Assertion: New Philosophical Essays* (pp. 277–296). Oxford University Press.
- Marmor, A. (2009). *Social Conventions: From Language to Law*. Princeton University Press.
- Marsili, N. (2019). The norm of assertion: A ‘constitutive’ rule? *Inquiry*, 0(0), 1–22. <https://doi.org/10.1080/0020174X.2019.1667868>
- McMyler, B. (2011). *Testimony, Trust, and Authority*. Oxford University Press.
- Mele, A. R. (1995). *Autonomous Agents: From Self Control to Autonomy*. Oxford University Press.
- Mele, A. R., & Moser, P. K. (1994). Intentional Action. *Noûs*, 28(1), 39–68. <https://doi.org/10.2307/2215919>
- Millikan, R. G. (1987). *Language, Thought, and Other Biological Categories: New Foundations for Realism* (Reprint edition). MIT Press.
- Moran, R. (2018a). *The Exchange of Words: Speech, Testimony, and Intersubjectivity*. Oup Usa.
- Moran, R. (2018b). *The Exchange of Words: Speech, Testimony, and Intersubjectivity*. Oxford University Press.
- Murray, S. E., & Starr, W. B. (2021). The Structure of Communicative Acts. *Linguistics and Philosophy*, 1–50. <https://doi.org/10.1007/s10988-019-09289-0>
- Neander, K. (1991). “Functions as Selected Effects: The Conceptual Analyst’s Defense.” *Philosophy of Science* 58(2): 168–84.
- Owens, D. (2006). “Testimony and Assertion.” *Philosophical Studies* 130(1): 105–29.
- Pagin, P. (n.d.). *Against Normative Accounts of Assertion*.
- Pagin, P. (2004). Is assertion social? *Journal of Pragmatics*, 36(5), 833–859. <https://doi.org/10.1016/j.pragma.2003.10.004>
- Pavese, C. (2018). Know-how, action, and luck. *Synthese*. <https://doi.org/10.1007/s11229-018-1823-7>

- Pavese, C. (2020). Knowledge, Action, and Defeasibility. In J. Brown & S. Mona (Eds.), *Reasons, Justification, and Defeaters*. Oxford University Press.
- Pettit, P. (1996). Functional Explanation and Virtual Selection. *British Journal for the Philosophy of Science*, 47(2), 291–302. <https://doi.org/10.1093/bjps/47.2.291>
- Pinker, S., Nowak, M. A., & Lee, J. J. (2008). The logic of indirect speech. *Proceedings of the National Academy of Sciences*, 105(3), 833–838. <https://doi.org/10.1073/pnas.0707192105>
- Pollack, M. E. (1992). The Uses Of Plans. *Artificial Intelligence*, 57, 43–68.
- Pritchard, D. (2020). Knowledge, skill and virtue epistemology. In E. Fridland & C. Pavese (Eds.), *Routledge Handbook of Skill and Expertise* (pp. 135–145). Routledge.
- Roberts, C. (2018). Speech Acts in Discourse Context. In *New Work on Speech Acts*. Oxford University Press.
- Searle, J. R. (1970). *Speech Acts: An Essay in the Philosophy of Language*.
- Shepherd, J. (2021). *The Shape of Agency: Control, Action, Skill, Knowledge*. Oxford: Oxford University Press.
- Sorensen, R. (2007). Bald-Faced Lies! Lying Without the Intent to Deceive. *Pacific Philosophical Quarterly*, 88(2), 251–264. <https://doi.org/10.1111/j.1468-0114.2007.00290.x>
- Stalnaker, R. (2006). Saying and Meaning, Cheap Talk and Credibility. In A. Benz, G. Jäger, & R. van Rooij (Eds.), *Game Theory and Pragmatics* (pp. 83–100). Palgrave Macmillan UK. https://doi.org/10.1057/9780230285897_2
- Stalnaker, R. (2014). *Context*. Oxford University Press.
- Stalnaker, R. C. (2019). *Knowledge and Conditionals: Essays on the Structure of Inquiry*. Oxford, England: Oxford University Press.
- Stanley, J. (2005). *Knowledge and Practical Interests*. Oxford University Press.
- Starr, W. B. (2018). *Socializing Pragmatics* (Vol. 1). Oxford University Press. <https://doi.org/10.1093/oso/9780198791492.003.0007>
- Turri, J. (2016). *Knowledge and the Norm of Assertion: An Essay in Philosophical Science* (1st ed.). Open Book Publishers. <https://www.jstor.org/stable/j.ctt1bpmb9w>
- Whiting, D. (2013). “Stick to the Facts: On the Norms of Assertion.” *Erkenntnis* 78(4): 847–67.
- Weiner, M. (2005). Must We Know What We Say? *The Philosophical Review*, 114(2), 227–251.
- Williamson, T. (1996). Knowing and Asserting. *The Philosophical Review*, 105(4), 489–523. JSTOR. <https://doi.org/10.2307/2998423>
- Williamson, T. (2000). Knowledge and its Limits. In *Knowledge and its Limits*. Oxford University Press.

Williamson, T. (2017). Acting on Knowledge. In J. A. Carter, E. C. Gordon, & B. W. Jarvis (Eds.), *Knowledge First: Approaches in Epistemology and Mind* (pp. 163–181). Oxford: Oxford University Press.