




UNIVERSITAT_{DE}
BARCELONA

Themes in linguistic understanding

Cognition and epistemology

Jędrzej Piotr (J.P.) Grodniewicz



Aquesta tesi doctoral està subjecta a la llicència **Reconeixement- NoComercial 4.0. Espanya de Creative Commons.**

Esta tesis doctoral está sujeta a la licencia **Reconocimiento - NoComercial 4.0. España de Creative Commons.**

This doctoral thesis is licensed under the **Creative Commons Attribution-NonCommercial 4.0. Spain License.**

Themes in linguistic understanding

Cognition and epistemology

Jędrzej Piotr (J.P.) Grodniewicz

A thesis submitted for the degree of Doctor of Philosophy at the
University of Barcelona

Supervisor: Manuel García-Carpintero

Co-supervisor: Josep Macià



UNIVERSITAT^{DE}
BARCELONA

University of Barcelona

Faculty of Philosophy

PhD Program: Cognitive Science and Language

(EEES: HDK08 Ciència Cognitiva i Llenguatge)

Line of Investigation: Analytic Philosophy (101133 Filosofia Analítica)

September 2020

Abstract

This thesis investigates the nature and epistemic role of linguistic understanding. It consists of five largely autonomous chapters. Chapter 1 focuses on the dynamic and temporal aspect of linguistic understanding. It argues that linguistic understanding can be appropriately characterized as a process and offers such a characterization. Chapter 2 argues that the process of language comprehension generates multiple interdependent representations (and metarepresentations). Only by taking into account all these representations and relations between them, we will be able to describe how linguistic understanding contributes to our acquisition of beliefs and knowledge. Chapter 3 argues that the justification of *comprehension-based beliefs*, i.e., the beliefs about what other people say, is non-inferential. In particular, it defends a version of process-reliabilism about the justification of comprehension-based beliefs. Chapter 4 examines the effectiveness of the filtering of information we acquire through comprehension. It argues that the filtering is not *real-time effective*, but it is *long-term effective*; it does not allow a particular hearer on a particular occasion to respond discriminately to a particular instance of testimony, but it shapes our social environment in such a way that untrustworthy testimony is relatively uncommon. As a result, the chapter supports a version of strong anti-reductionism about testimonial entitlement. Finally, Chapter 5 argues that linguistic understanding differs from the so-called *understanding of a proposition*. The latter is a kind of what contemporary epistemologists characterize as objectual understanding. Nevertheless,

in most cases of successful linguistic communication, linguistic understanding, and understanding of a proposition jointly contribute to *understanding a communicated thought*.

Contents

Acknowledgments	iv
Introduction	vi
1 The process of linguistic understanding	1
2 The representational structure of linguistic understanding	27
3 The justification of comprehension-based beliefs	54
4 Effective filtering	78
5 Understanding a communicated thought	104
Summary	124
References	127

Acknowledgments

I wouldn't be able to write this thesis without support from a group of wonderful people. First, I would like to thank my main supervisor Manolo García-Carpintero. Never imposing anything, Manolo was always available with his helpful guidance and insightful suggestions. He also had limitless patience for my attempts to tackle research problems from the most unexpected angles and the trust that I would eventually find my way. It was very encouraging. Thank you, Manolo! I would also like to thank my second supervisor Josep Macià. Josep is definitely one of the most meticulous readers I know. Our discussions were always challenging and extremely thought-provoking. Thank you, Josep!

LOGOS research group at the University of Barcelona is a perfect environment to forge one's academic skills and grow as a philosopher. As a PhD student at LOGOS, I have benefitted from hundreds of hours of discussions and seminars with LOGOS members and visitors. Special thanks to Anton Alexandrov, Andre Bazzoni, Samuele Chilovi, Filippo Contessi, Diana Couto, Roger Deulofeu Batllori, Esa Díaz-León, Johannes Findl, Alfonso García Lapeña, Ramón García Moya, Carlos Jaén Solanes, Max Kölbel, Alex Krauss, David J. Lobina, Teresa Marques, Neri Marsili, Manolo Martínez, Sergi Oms, Michele Palmira, Andrea Rivadulla Duró, Adam Sennet, Javier Suárez, Josefa Toribio, Matheus Valente, and Lisa Vogt.

During my study visits, I was lucky to have two amazing sponsors: Liz Camp (at Rutgers University) and Derek Ball (at the University of St An-

draws). Both of them were extremely generous with their time and expertise. This thesis benefitted greatly from their input. I am also very grateful to Carolina Flores and Katharina Bernhard for being such excellent guides through the academic and social life of Rutgers and St Andrews, respectively. Finally, my stay at St Andrews would have been much gloomier, if Christophe Geudens, with whom I have shared my office, wouldn't be so eager to take all the breaks we took, and drink all the coffees we drank.

Many thanks to J. Adam Carter and Emma C. Gordon for the inspiring and efficient collaboration on the paper "Understanding a communicated thought" (Chapter 5 of this thesis).

The material for this thesis has been presented at conferences and workshops in Barcelona, Cambridge, Łódź, Herdfordshire, Munich, San Sebastian, Warsaw, Rijeka, Vienna, Kraków, Oslo, Glasgow, and St Andrews. In all these places I have received plenty of valuable feedback. Special thanks to Tadeusz Ciecierski, Chris Gauker, Peter J. Graham, Leopold Hess, Mark Jary, Natalia Karczewska, Christoph Kelp, Palle Leth, Olga Poller, Krzysztof Poślajko, and Mona Simion.

Bearing all the challenges of academic life (and life in general!) has been much easier thanks to my friends: Andrzej Bułeczka, Bartek Czajka, Grzegorz Gaszczyk, Aleksandra Krogulska, Tomasz Pietrzak, Abel Suñé, Anna Wilczyńska, Karol Wilczyński, and Anna Żymełka-Pietrzak. Moreover, Bartek and Grzegorz have spent dozens and dozens of hours reading and commenting on my work. Thank you!

Finally, I wouldn't even consider undertaking challenges like a PhD in philosophy without the drive and self-confidence instilled and sustained in me by my parents and my brother. I dedicate this work to them.

Introduction

Przytłaczająca obfitość, powiązań, skojarzeń... Ileż zdań można utworzyć z dwudziestu czterech liter alfabetu? Ileż znaczeń można wyprowadzić z setek chwastów, grudek i innych drobiazgów?

An overwhelming abundance of connections, associations... How many sentences can one create out of the twenty-four letters of the alphabet? How many meanings can one glean from hundreds of weeds, clods of dirt, and other trifles?

— Witold Gombrowicz, *Cosmos*¹

If we examine the traditional borders of the philosophical sub-disciplines, we readily realize that linguistic understanding is, to some extent, *terra nullius* — nobody's land. It lies no more in the domain of the philosophy of language than the philosophy of mind. It is hard to imagine anyone safely crossing its treacherous paths without the equipment provided by epistemology and cognitive sciences. And yet, it is a central and strategically important territory. Whoever investigates human social activity will, sooner or later, arrive at its borders. Will they have maps precise enough to navigate it without distress?

¹English translation by Danuta Borchardt.

This thesis is an attempt to fill in some of the blank spots on our maps of linguistic understanding. It consists of five essays. Each of them can be considered a separate expedition to a specific problem area. Each required a slightly different toolkit and preparation. Nevertheless, all the explored territories adhere to one another, and everything that I have discovered along my way allowed me to later move more confidently into the yet unknown.

Chapter 1 starts with a simple but significant observation: spoken linguistic communication is very fast. Right after one person finishes their utterance, another starts theirs. Is there even time to understand each other? Fast communicative *turn-taking* — as this phenomenon is called in conversation analysis and psycholinguistics — draws our attention to the dynamic and temporal aspect of linguistic understanding. I argue that we will not be able to accommodate this data if we focus exclusively on understanding characterized either as a *state* of grasping the content of a given utterance or as a *disposition* to reach these states, i.e., linguistic competence. Drawing from both metaphysics of temporal entities and psycholinguistics, I develop an account of the process of linguistic understanding. I argue that linguistic understanding, just like other processes, unfolds over time. Understanding of an utterance begins right after we hear the first sounds of this utterance and continues until the final state of understanding of this utterance is produced.

In Chapter 1, I also propose the following stipulative definition combining all aspects of linguistic understanding:

(LU) Linguistic understanding is a cognitive process that typically takes place on hearing an utterance in a language one disposition-understands and which typically results in a set of mental states via which the hearer represents the content and force of the utterance or what the hearer identifies as the content and force thereof.

Chapter 2 focuses on the products of the process of understanding, i.e., states of understanding. It starts with a comparison of the five most influen-

tial philosophical theories of states of understanding. Three of them, I call them *indirect*, characterize understanding as a state which meta-represents the content of an utterance *as said by a given speaker*, e.g., *that Anne said that p*. Two others, I call them *direct*, characterize understanding as a state which represents just the content of the utterance (*p*). I argue that both these approaches are only partially right. Instead of developing another single-state characterization of linguistic understanding, I offer a model of the representational structure of linguistic understanding consisting of multiple interdependent representations (some of which are direct, some indirect).

In this model, language comprehension runs in two streams: faster, which contains only direct representations of the content of an utterance, and slower, which contains metarepresentations of the content *as said by a given speaker*. Each of these streams is equipped with a dedicated filter. The representations in the faster stream are filtered by *validation*, i.e., a purely content-oriented filter, which checks comprehended contents for consistency with our active or easily accessible background knowledge. The representations in the slower stream, on the other hand, pass through a source-oriented filter, which monitors whether the comprehended information comes from a reliable source.

Most importantly, the model I offer in Chapter 2, establishes a middle ground between the so-called *Spinozan* model of comprehension (according to which we, by default, believe everything we comprehend), and the so-called *Cartesian* model of comprehension (according to which we first entertain the comprehended content and then either accept or reject it). I argue that we do not accept *everything* we comprehend, but we accept everything that passes validation, i.e., the filter localized in the faster stream. This conclusion is crucial for my discussion of the effective filtering of testimony in Chapter 4.

In **Chapter 3**, I take a closer look at the so-called *comprehension-based beliefs*, i.e., the beliefs whose contents are meta-representations of the form *that the speaker said that p*. The question is: how are these beliefs justified?

According to *epistemic non-inferentialism*, the justification of these beliefs does not depend on the justification of any other beliefs. The most influential version of epistemic non-inferentialism is *language comprehension dogmatism* — a view that comprehension-based beliefs are *prima facie* justified by experiential states, i.e., seemings or quasi-perceptions *that the speaker said that p* (Fricker, 2003; Brogaard, 2018). Recently, this position has been attacked by Brendan Balcerak Jackson (2019), who provided three counterarguments against language comprehension dogmatism. Moreover, Balcerak Jackson took these arguments to constitute an abductive argument for *epistemic inferentialism* — a view according to which comprehension-based beliefs are justified by other beliefs, e.g., beliefs about what words the speaker uttered.

I argue that Balcerak Jackson’s arguments may be effective against language comprehension dogmatism (even though they are much less damaging than he thinks), but do not constitute a successful abductive argument for epistemic inferentialism. I develop another non-inferentialist position, i.e., *comprehension-process reliabilism*, which deals with Balcerak Jackson’s arguments.

Chapter 4 goes one step further into the territory of the epistemic role of language comprehension. It takes up the question of whether our ability to filter out unreliable testimony is effective. My answer has two parts. First, I argue that filtering of testimony is not *real-time effective*, i.e., we are not able to effectively detect untrustworthy testimony while (or right after) comprehending it. This claim, already defended by Michaelian (2010) and Shieber (2012), has been recently contested by Fricker (2016). I argue, not only that Fricker’s critique of Michaelian and Shieber fails, but also that real-time effective filtering is impossible if the model of language comprehension, which I offered in Chapter 2, is correct.

Second, I argue that filtering is *long-term effective*. The fact that all, at least all adult members of our linguistic communities are routinely vigilant (i.e., monitor incoming testimony) and that being caught on a lie triggers

social retribution, puts pressure on speakers who are then much less inclined to produce false testimony. Therefore, filtering is effective not as an ability to detect falsehoods on the spot, but as a force shaping our social reality. In the last section of Chapter 4, I argue that these considerations support a version of strong anti-reductionism about testimonial entitlement, i.e., a view according to which we are *prima facie* entitled to believe whatever we are being told.

The final chapter (**Chapter 5**) is another exploration of almost virgin land. It examines the phenomenon of what I and coauthors of this chapter: J. Adam Carter and Emma C. Gordon call *understanding of a proposition*. We argue that understanding of a proposition (often identified with linguistic understanding) is a kind of what contemporary epistemology characterizes as *objectual understanding*. Even though linguistic understanding and understanding of a proposition are distinct phenomena, in everyday communication, they often jointly contribute to *understanding a communicated thought*. To understand a communicated thought, one has to, first, recognize which content has been expressed in a given utterance (linguistic understanding), and second, understand this content itself (understanding of a proposition).

Chapter 1

The process of linguistic understanding¹

Abstract. The majority of our linguistic exchanges, such as everyday conversations, are divided into turns; one party usually talks at a time, with only relatively rare occurrences of brief overlaps in which there are two (or more) simultaneous speakers. Moreover, conversational turn-taking tends to be very fast. We typically start producing our responses before the previous turn has finished, i.e., before we are confronted with the full content of our interlocutor's utterance. This raises interesting questions about the nature of linguistic understanding. Philosophical theories typically focus on linguistic understanding characterized either as an ability to grasp the contents of utterances in a given language or as outputs of this ability—mental states of one type or another. In this chapter, I supplement these theories by developing an account of the process of understanding. I argue that it enables us to capture the dynamic and temporal aspect of understanding

¹Material from: Grodniewicz, J.P., "The process of linguistic understanding." *Synthese* (2020), doi.org/10.1007/s11229-020-02807-9, Springer Nature.

and reconcile philosophical investigations with empirical research on language comprehension.

1 Introduction

According to conversation analysis studies, the majority of our linguistic exchanges, such as everyday conversations, are divided into turns (Sacks et al., 1974). That is to say that one party usually talks at a time, with only relatively rare occurrences of brief overlaps in which there are two (or more) simultaneous speakers. This observation would not be striking if it were not juxtaposed with another empirical result provided by psycholinguistics: the gaps between subsequent turns are very short, only about 0.2 s (Stivers et al., 2009). Given that language production requires a minimum of around 1 s for an unprimed single-word utterance to around 1.5 s for a simple clause, we have to conclude that subsequent turns are planned, and that the production process begins before the end of the previous turn (Levinson and Torreira, 2015; Garrod and Pickering, 2015; Levinson, 2016). Consequently, speakers typically start responding not to the whole content of an utterance they have heard, but to their prediction of the full content of the utterance on the basis of only that part of the content they have actually heard.² To account for these data, we need a theory of linguistic understanding more versatile and of a broader scope than those already available in the philosophical debate.³

²Of course, to say that 0.2 s is the length of a typical gap in conversation is a vast generalization. A lot depends on such factors as the type of the upcoming turn. For example, Kendrick (2015) observes that turns involving repair appear approximately 0.7 s after the preceding ones. However, even 0.7 s is still a short enough gap to support the argument presented in what follows.

³It is important to keep in mind that linguistic understanding that is the subject of this chapter should not be identified with communicative success. Arguably, communicative success may, in specific situation, be achieved despite the lack of linguistic understanding (or in the case of only partial understanding). Thus, linguistic understanding is not identical with the communicative success, although might (and typically does) contribute to it.

Philosophical theories typically focus on states of understanding, i.e., “outputs of exercises of one’s ability” to understand (Longworth, 2018). They look for an answer to a very specific question: What is the hearer’s attitude towards the content or *what is said* through an utterance which that hearer understands? Most influential theories characterize understanding either as a kind of: knowledge (e.g., Evans, 1982; Dummett, 1993; Davies, 1989; Heck, 1995; Campbell, 1982; Higginbotham, 1992), belief (e.g., Millikan, 1984, 2004), perception-like state of direct awareness of the content (and force) of the utterance (Hunter, 1998; Fricker, 2003), or a state through which the hearer entertains the content of the utterance (Longworth, 2018).

Research on turn-taking shows that we start producing our responses before the turn in progress at the time has finished, i.e., before we are confronted with the full content of our interlocutor’s utterance. Certainly, it is quite natural to assume that a hearer is in a state of understanding an utterance as soon as the last word of that utterance has been heard. However, short turn-taking gaps simply do not provide enough time for a hearer to produce a response only when already in a state representing the content of the previous turn. Thus, apparently, in the vast majority of communicative situations we produce our conversational turn without being in one of the states typically characterized by philosophers as the state of understanding (I come back to this matter in Section 4). In result, to account for the role it plays in fast communicative exchanges, we have to extend our characterization of linguistic understanding beyond the most popular state-sense. The authors enumerated above do not claim that being in a state of understanding is *all* there is to linguistic understanding. However, none of them provide an elaborated account of the dynamic aspect of linguistic understanding. This chapter is meant as an attempt to provide such an account.⁴

It is important to keep in mind that the current chapter is not meant as

⁴In this essay, I focus on the understanding of speech (spoken linguistic input). I think that it will be relatively easy to extend my proposed account to the written word, although I do not undertake this task here.

a critique of the available theories of linguistic understanding listed above. The states of understanding that they focus on are independently interesting. We want such states to play various roles in our cognitive enterprises. For example, states of understanding figure as partial reasons for action (e.g., I have run away because I knew that what you have said was that there is a tiger behind the rock), and belief formation (I believe that there is a tiger behind the rock because this is what you have told me). Some theory of states of understanding has to be a part of a more versatile theory of linguistic understanding I am about to offer.⁵

My plan is the following. In Section 2, I present three senses of linguistic understanding. In Section 3, I discuss the relation between these senses and develop an account of the process of understanding. In Section 4, I respond to a possible objection according to which we can account for the dynamic aspect of linguistic understanding without characterizing linguistic understanding as a process. Finally, in Section 5, I outline additional benefits of my account for the philosophical debate about linguistic understanding.

2 Three senses of “linguistic understanding”

There are at least three senses in which we pre-theoretically grasp the idea of linguistic understanding (Hunter, 1998; Longworth, 2010). I have already mentioned the first: we often think of understanding as a mental state in which a subject represents the content of a heard utterance. In the second sense, we consider understanding to be a disposition or stable ability to decode sentences of a certain language (I will be sometimes referring to it as *disposition-understanding* of a given language). This is what we mean when we talk of understanding (or simply “speaking”) a given language, such as English, Kashmiri, or Northern Sotho.

⁵I outline my view on states of understanding in Chapter 2. Moreover, in Section 4, I discuss the hypothesis that the process of understanding is constituted out of states of (partial) understanding.

It is worth noting that we expect states of understanding (the first sense of understanding given above) to be grounded in a disposition (this second sense of understanding). Fricker (2003) highlights the fact that to have an understanding of a particular utterance of a sentence, *S*, is to enjoy a correct representation of the content of the utterance “as a result of one’s dispositional understanding *S*—that is, through the exercise of a stably possessed, internally constituted reliable capacity” (p. 346). It is not enough to simply grasp the content of an utterance. As Fricker shows through her example, I might effectively grasp the contents of Russian utterances if there is a reliable translator by my side who translates from Russian to English for me. Does that mean that I understand Russian? In a sense, I do; in a sense, I do not. If, asking about understanding, we ask whether I can communicate with my Russian monolingual friend, it may seem that I understand Russian (thanks to my Russian translator!). However, we might be concerned with a more profound matter: whether I myself as a speaker with such and such biological cognitive setting, such-and-such experience and language acquisition history understand Russian. In this latter case the answer would be “no”. It is obvious that what is of primary importance for our research into the nature of linguistic understanding is the latter question, and the sense in which I do not understand Russian. At least when trying to refine the basic sense of linguistic understanding (which could, later on, be used to set up standards for possible derivative senses such as “I don’t know any Russian but I could understand him because I had Google Translate on my phone.”), we need to focus on a disposition to understand a language which is common to almost all human beings and which is internalized in a very specific way. It is grounded in our language faculty (e.g., Hauser et al., 2002; Jackendoff and Pinker, 2005) consisting, among others, of an appropriate biological setting of a language-ready brain (e.g., Boeckx and Benítez-Burraco, 2014; Friederici, 2012), some form of internalization of the grammar of a given language (e.g., Pinker and Bloom, 1990), a mental lexicon (e.g., Jackendoff, 2002; Patterson

et al., 2007), etc. Due to the lack of space, I will leave a discussion of the exact nature of that disposition for another occasion. However, I agree with Fricker that possession of such disposition is necessary for being in states of understanding of particular utterances. States of understanding of utterances of sentences in a given language result, at least in part, from exercising the disposition to understand this language.

The third way in which we can think of understanding is to think of it as an “event” (Hunter, 1998, 561), “process” (Fricker, 2003, 330), “episode” or “achievement” (Longworth, 2010, 4) of one’s coming to the state of understanding of a given utterance.⁶ Notably, there is relatively little discussion about this sense of understanding, either in the philosophy of language or the philosophy of mind. Someone might even assume, that the nature of processes, events, or episodes of understanding lies outside the area of philosophical considerations; it is the task of psycholinguistics to study time-spans and the mechanisms of language comprehension, and that of neurolinguistics to study the neural pathways recruited for the processes, etc. However, as I argue in the next section, focusing on the dynamic and temporal character of linguistic understanding opens a rich and interesting avenue of philosophical research.

3 Linguistic understanding as a cognitive process

3.1 The metaphysics of processes

How do the three senses of linguistic understanding mentioned above hang together? Are they all pointing at entirely distinct phenomena or rather different aspects of the same complex phenomenon? Finally, is it possible that careful examination of any one of these senses will help us to better

⁶I will discuss the differences between these notions in the next section.

understand the remaining two?

I think that the three senses of linguistic understanding indicate different aspects of the same phenomenon, and as such, they are strongly interdependent. I will now propose a stipulative definition of linguistic understanding, which outlines the relations between them. Given that my primary interest in this chapter is the dynamic and temporal aspect of linguistic understanding, my definition puts the process sense of linguistic understanding to the foreground. Crucially, the definition is, at least *prima facie*, compatible with a variety of views about the nature of state-understanding and disposition-understanding available in the debate.

(LU) Linguistic understanding is a cognitive process that typically takes place on hearing an utterance in a language one disposition-understands and which typically results in a set of mental states *via* which the hearer represents the content and force of the utterance or what the hearer identifies as the content and force thereof.

LU assumes very little about the nature of the process of understanding. In principle, it is possible that, e.g., Cyborgs or Martians were capable of performing *the same* process, even if they were equipped with different cognitive and perceptual mechanisms than these possessed by humans. Independent of a specific implementation, a process that is based on a disposition to understand a given language and produces mental states representing contents of utterances of sentences in this language, qualifies as a process of understanding.

Moreover, LU specifies what triggers the process of linguistic understanding. Just as the process of watching a particular material object requires a particular material object that is being watched (Crowther, 2009), the process of linguistic understanding requires some linguistic input that is being understood. However, restriction to *typical* cases is important here. Not

always when a subject possessing a disposition-understanding of a given language encounters utterances in this language, the process of understanding is being triggered. If the hearer is cognitively impaired or if their cognitive system has been tinkered with or, finally, if their attention resources are insufficient it is possible that they would encounter an utterance but the process would not be triggered or the process would be triggered but the state would not be produced (more on this below). As is the case with other types of processes dependent on the working of our cognitive mechanisms, we should restrict our predictions about linguistic processing to broadly construed *normal conditions*.

But is it appropriate to think about linguistic understanding as a process? To answer this question, I will draw from a rich tradition of philosophers who investigated the nature of temporal entities, often in the context of the discussion about grammatical aspect (cf. Vendler, 1957; Verkuyl, 1972; Mourelatos, 1978; O’Shaughnessy, 2000; Rothstein, 2004; Crowther, 2009, 2011, 2018; Steward, 1997, 2012, 2018; Galton, 2018; Soteriou, 2013, 2018; Stout, 2018). Different authors use different classifications but I will distinguish four basic categories: states, achievements, accomplishments, and activities. My goal in this section is to argue that there are senses of “linguistic understanding” belonging to each of these categories. Is it even possible? It is. As indicated by Vendler:

[it is not the case that] a verb exhibiting a use fairly covered by one schema cannot have divergent uses, which in turn may be described in terms of other schemata. As a matter of fact, precisely those verbs that call for two or more time schemata will provide the most interesting instances of conceptual divergence in this respect—an ambiguity which, if undetected, might lead to confusion. Thus my intention is not to give rules about how to use certain terms but to suggest a way of describing the use of those terms. (Vendler, 1957, 143-4)

As I have already mentioned, there is an obvious sense in which “linguistic understanding” denotes a *state*. Depending on which theory we favor, we can characterize the state of understanding of an utterance either as knowledge that what was said was so and so, a belief that what was said was so and so, a state of entertaining the content of this utterance, etc. Crucially, “states such as knowledge and belief do not have temporal duration in the sense of having temporal parts or successive temporal phases over which they unfold.” (Crowther, 2011, 5). The same state of my understanding of your utterance that p may obtain now and in an hour, and it does not make sense to say that in an hour I am in a different phase, or at a different stage of this state.

Importantly, disposition-understanding which we pick out when we say “Anne understands Italian” is also a state, although different than a state of understanding of a particular utterance. As I said above, disposition-understanding of a given language is a state of a subject’s cognitive system, which enables the system to respond to linguistic input, e.g., particular utterances, by producing states of understanding of these utterances.

The next category enumerated above is this of achievements. Achievements are typically characterized as instantaneous changes of state. Some clear examples are *noticing*, *recognizing*, *finding*, and *dying*. As observed by Longworth (2010), there is an intuitive sense in which *linguistic understanding* denotes an achievement. In this sense, understanding is an episode of: “one’s coming to state-understand, or understanding, what someone has said” (Longworth, 2010, 4); it is the temporal boundary between whatever led to the production of state-understanding and the first moment in which the state-understanding obtains.

Linguistic understanding has been previously characterized as a state and as an achievement, so these two senses are fairly uncontroversial. But I have promised to argue that there are senses of linguistic understanding corresponding to all four categories enumerated above, that is also to activities and accomplishments. In what follows, I will use the term “process” as a

collective name for activities and accomplishments.⁷ Therefore, my goal now is to argue that we can think about linguistic understanding as a process.

Activities and accomplishments are both progressive, i.e., they describe temporal entities unfolding over time. Names of particular activities and accomplishments can be appropriately used as answers to questions such as “What are you doing?” or “What is going on?” (Steward, 2018, 111). The core difference between activities and accomplishments concerns telicity.⁸ Accomplishments such as *walking to the shops* or *writing a letter* are oriented towards specific goals or points “at which they must terminate” (Crowther, 2011, 6), while activities, such as *walking* or *writing*, are not. If I have walked for ten minutes and after that kept walking for another five minutes, there was an activity of walking that took fifteen minutes. However, if I have walked to the shops in ten minutes, it means that I have reached the shops after ten minutes. After that, I can walk to the shops *again*, but I cannot walk to the shops *some more*.⁹

Understanding an utterance is a candidate for an accomplishment, because it has specific “completion condition” (Crowther, 2011), namely, the achievement of the state of understanding. The claim I want to defend is that *understanding an utterance* unfolds over time because *hearing an utterance* unfolds over time and the temporal structure of understanding an utterance overlaps with the temporal structure of hearing this utterance.¹⁰ In Section 3.2, I present psycholinguistic data about incrementality and predictivity of language understanding, which support this claim. Here, my focus is on its philosophical motivation.

⁷Other authors pick different terminological conventions. For example, Crowther (2011) uses the term “process” for what I, following e.g., Vendler (1957) and Rothstein (2004), call “activity.”

⁸Cf. Rothstein (2004).

⁹For an exhaustive and illuminating discussion of the difference between activities and accomplishments see (Crowther, 2011).

¹⁰That does not mean that they have *the same* temporal structure. More on this topic in Section 4.

Let us consider the following, particularly interesting example of an accomplishment provided by Matthew Soteriou (2018):

It takes around 51 minutes and 14 seconds to play, from start to finish, the 1981 recording of Glenn Gould’s performance of the Goldberg Variations. It takes at least the same amount of time to hear it. Once you’ve listened to the whole recording, you will have heard the final recording of Gould’s performance of that work. . . . Although we’re unable to hear a performance of the complete work all at once, so to speak, this doesn’t mean that we are incapable of hearing a performance of the whole work. It just means that this is an accomplishment that takes time. (Soteriou, 2018, 82)

The reason why it takes 51 minutes and 14 seconds to hear the 1981 recording of Goldberg Variations is that playing this piece takes the same amount of time. At each moment of playing, there is simultaneous hearing. But, let’s go one step further and imagine that the hearer in Soteriou’s example is a music critic whose goal is to analyze the 1981 recording. Most likely, the critic does not wait with their analysis till the *playing* is finished and the *hearing* accomplished. They analyze Goldberg Variations as they hear them being played. So, it takes (at least) 51 minutes and 14 seconds to analyze the 1981 recording of Goldberg Variation.

I submit that, in this respect, *understanding an utterance* is analogous to *analyzing the 1981 recording of Goldberg Variations*.¹¹ If you utter the sentence “It would be really nice to visit Vienna again, like in the old times,” the uttering of this sentence takes time. Not as much time as playing the Goldberg Variations, but still, it unfolds from the first word to the last over

¹¹Certainly, analyzing the *recording of Goldberg Variations* may continue long after the *playing* and *hearing* are finished (in fact, it may take place over the course of multiple *playings* and *hearings*), while *understanding an utterance* typically finishes right after the *vocalizing* and *hearing*.

the course of 3 seconds or so. My hearing of this utterance, analogously to Soteriou's example, takes more or less the same amount of time as it takes vocalizing it. Crucially, my understanding of this utterance starts when I hear the first word and continues as I hear subsequent words. Just as the critic does not wait with analyzing the Goldberg Variations till the accomplishments of *playing* and *hearing the Goldberg Variations* are finished, our cognitive systems do not "wait" with understanding an utterance till the end of accomplishments of *uttering* and *hearing this utterance*.

Obviously, the accomplishment *understanding an utterance* is intimately related with both state of understanding and an achievement of understanding. The state of understanding an utterance is the *telos* towards which the accomplishment progresses. The achievement, as indicated by Longworth, is the boundary between the accomplishment and the first moment of being in the state of understanding; an instantaneous episode of coming to state-understanding. Stout observes:

In English the form of words used to refer to an achievement is very often the same as the form of words used to refer to a process. The phrase 'my crossing the road' serves both purposes. This is an instance of the familiar process/product ambiguity. We might say that the writing that was happening in the exam hall resulted in the writing that could be found later on the exam papers. The word 'writing' is functioning differently in its two occurrences here, first to refer to a process and then to refer to the product of that process. (Stout, 2018, 212)

Arguably, in contemporary English, "understanding" is almost never used as progressive. We use expressions like "Anne is crossing the road." or "Anne is analyzing the data." but "Anne is understanding the data." does not sound right. In this context, it is much more natural to say, e.g., "Anne is trying to understand the data." Is it a problem for my project of characterizing linguistic understanding as an accomplishment?

I do not think it is. Even if the progressive use of “understand” is not common, it is not unprecedented.¹² Moreover, the frequency and context in which different words are used in the progressive undergo diachronic changes and are subject to dialectical differences (Smutterberg, 2005; Levin, 2013). We encounter at least a trace of progressiveness in constructions like: “If I’m understanding the results correctly...” This phrase can be interpreted as conveying both that I am in a state of understanding the results in some way (and I am not sure whether this way is the right way) and that I am still in the process of understanding these results (which, hopefully, will culminate with understanding them correctly). To make the process sense even more explicit, imagine that I am reading the results as they are being printed page-by-page. After reading just a few pages, I observe certain tendencies and immediately report: “If I’m understanding the results correctly, the asteroid is going to hit us in 16 hours.” I am still in the process of understanding these results, but my current best guess is that the results indicate that the asteroid will hit us in 16 hours. I take the process of linguistic understanding to be analogous to the process of understanding the results as they are being printed page-by-page. Again, following Vendler, my goal is “not to give rules about how to use certain terms but to suggest a way of describing the use of those terms.” (Vendler, 1957, 144).

But would it not be better to simply drop the controversial progressive use of “understand” and use *coming to understand an utterance* when speaking about accomplishments? The problem with using *coming to understand* is that it suggests that understanding “is still entirely in the future” (Galton, 2018, 55). However, as my comparison with *analyzing the Goldberg Variations* was intended to demonstrate, and as I will further argue by appeal to empirical research in the next section, when we listen to someone vocalizing an utterance, understanding is not entirely in the future. It is in

¹²See Jørgensen (1991) for a discussion of progressive uses of predominantly stative verbs including “understand.”

progress from the get-go and unfolds over time, simultaneously with *vocalizing* and *hearing*. Word after word there is more and more *understanding an utterance* (an accomplishment) leading to *understanding an utterance* (an achievement), which, in turn, demarcates the first moment of *understanding an utterance* (a state). Vendler is on point: “verbs that call for two or more time schemata will provide the most interesting instances of conceptual divergence.” (1957, 144).¹³

Finally, can we go one step further and think about the use of *linguistic understanding* as an activity? I think we can. Here is Rob Boffard describing his experience with a speed-reading app in a Guardian article: “I started off at a relatively mild 350wpm, but soon graduated to a cruising altitude of 600wpm. And I was understanding all of it.” (Boffard, 2014, September 22). The more *reading* (activity), the more *understanding* (activity). After reading (and understanding what he was reading) for ten minutes, Boffard could have read (and understood what he was reading) for another five minutes, and there would be fifteen minutes of reading with understanding, just like we expect from activities.

In this section I described various uses of *linguistic understanding*. Using *linguistic understanding* as denoting a state or an achievement, corresponding to the ways in which the verb “understand” is typically used in contemporary English, was already common in the debate. However, I argued that we can also use *understanding* as denoting a process (an accomplishment or an activity). From the philosophical point of view, there are good reasons to do so. The temporal structure of the process of understanding overlaps with

¹³Notice also, that (in the accomplishment sense) the sentence “Anne was understanding what Tom was saying” is consistent with its never being the case that she understood what Tom was saying. If Tom was saying “The results clearly indicate that the asteroid is going to hit us in 16 hours.” but he was wrong, and the asteroid hit right after he vocalized the word “asteroid” annihilating all life on earth, Anne was understanding what Tom was saying but she never understood it. We can say that she understood a part of it, but we can say the same thing about Alice who, when the asteroid hit, was *building a house* (a classical example of an accomplishment). Alice never built the house (achievement), even though she built a part of it, for example, just the walls.

temporal structures of other processes, both activities (e.g., *reading*) and accomplishments (e.g., *hearing an utterance*).

3.2 The process of understanding in psycholinguistics

Incrementality

Above, I have argued that, given its temporal structure, there are good reasons to think about linguistic understanding as a process. In this section, I will try to demonstrate that it perfectly fits the way in which *linguistic understanding* is being characterized by empirical language sciences.

Psycholinguistics commonly characterizes linguistic understanding as unfolding over time. A specific term has even been applied to the type of progressiveness in language comprehension. Psycholinguists say that language comprehension is “incremental,” i.e., new unfolding input is processed by the receiver chunk by chunk. This chimes well with the turn-taking data. Given the very short gaps between turns, it is simply impossible for a receiver’s cognitive system to “wait” until the speaker’s turn has finished and only then to “trigger” the comprehension process and produce a response. However, there are also other phenomena that support the claim regarding the incrementality of language comprehension. One of the most robust and widely studied is “garden pathing”, first identified by Thomas Bever (1970). There is a large and fast-growing body of literature on garden-path processing of written text comprehension; however, the phenomenon has also been extensively studied in the context of speech processing (Frazier and Fodor, 1978; Kjelgaard and Speer, 1999; Steinhauer et al., 1999; Schafer et al., 2000; Pauker et al., 2011). Take the sentence “The old man the boat.”. When receivers hear or read the full sentence, they realize that the most likely interpretation of the sentence’s meaning they had built after hearing or reading only “The old man” (i.e., that “the old man” is a noun phrase) is incorrect. As might be expected, this effect is blocked or reduced by appropriate prosodic markers (in speech) and

punctuation (in writing). Nevertheless, the very fact that the phenomenon occurs in some utterances is sufficient to show that language comprehension is incremental. If no interpretation had been formed before the full utterance was heard, receivers would not be surprised by how the utterance actually unfolds.

Another important argument in favor of the incrementality of language understanding comes from cross-modal eye-tracking experiments (Altmann and Kamide, 1999; Kaiser and Trueswell, 2004; Tanenhaus et al., 1995). In their influential study, Altmann and Kamide (1999) recorded the eye movements of participants presented with pictures while simultaneously listening to recorded sentences. The pictures were of objects with clearly differentiated properties, e.g., a boy, a cake, a car, a ball, and a toy train. The participants heard two types of sentences, e.g., “The boy will move the cake.” and “The boy will eat the cake.” Note that in the second sentence, only one of the objects from the example picture could be referred to as the object the boy was going to eat to form an utterance compatible with common sense (only the cake is eatable). As expected, when there was only one candidate object of a given verb, as in the second sentence above, the probability of the participants looking at that object was significantly increased. Altmann and Kamide conclude that: “information extracted at the verb can be used to guide eye movements to whichever object in the visual context satisfies the selectional requirement of the verb. This guidance is initiated before the linguistic expression corresponding to the verb’s direct object is encountered” (p. 258).

One could cite other similar phenomena, but I will follow the psycholinguistic mainstream in taking garden paths, turn-taking, and cross-modal eye tracking to constitute a sufficient argument for the incrementality of language comprehension. Just as suggested above, understanding does not happen after *hearing* or *reading* but happens simultaneously with them.

Prediction

There is another feature of linguistic understanding which speaks to its processual character. It is commonly assumed in psycholinguistics that understanding is not only incremental but also predictive. Understanding an utterance consists of constant formation and updating of predictions about what comes next; how the utterance will unfold and how it will end. Kuperberg and Jaeger (2016) go as far as to state that it is “logically impossible” to reject the claim that language comprehension is predictive in light of some robust linguistic phenomena. Two of those, garden-paths and turn-taking, I have already discussed.¹⁴ Additional evidence is provided by e.g., studies of eye movements during the processing of written texts. It has been repeatedly shown (see references in Kuperberg and Jaeger (2016)) that readers fixate more on unpredictable words and less on words that fit in well with the context (i.e., are easier to predict). For example, in a study conducted by Rayner and Well (1996), participants read sentences containing highly, moderately and weakly constrained target words while their eye movements were recorded. The sentences were taken from an earlier study by Schwanenflugel (1986), who asked participants to list up to three possible ways of completing a given sentence. For instance, for the beginning of a sentence: “The woman took the warm cake out of the . . .”, the word “oven” was produced more than ninety percent of the time; the word “stove”, slightly more than thirty percent; and the word “pantry” less than five percent of the time. Therefore, “oven” is highly constrained, “stove” is moderately constrained, and “pantry” weakly constrained completion of the target sentence. Rayner and Well observed that even though there were no significant differences in gaze-fixation

¹⁴Notice that the garden-path data is complementary to the turn-taking data. Fast turn-taking alone could result from interlocutors being simply uninterested in what others have to say and merely waiting for their chance to speak. Garden path data, however, proves that hearers build interpretations of what speakers say based on fragments of utterances which they have already perceived. This is why they are surprised when the actual continuations of utterances do not meet their predictions.

times while reading highly and moderately constrained target words, participants fixated for significantly longer on weakly constrained target words. This indicates that readers constantly predict what word or words will most likely appear next in a given context. Prolonged eye fixation on the unpredictable words reveals the increase in processing difficulty for unpredictable completions.

Let us briefly compare this with the asteroid case discussed in Section 3.1. When I reported that the asteroid is likely to hit us in 16 hours, I was still in the process of understanding the results (subsequent pages containing these results were still being printed), but my best guess about the results was that they indicate that the asteroid is going to hit us in 16 hours. This is similar to what we do when processing linguistic input. Imagine that I answer, “I’d like salad” 0.2 s after hearing “What would you like to have for lunch?” To answer so fast, I must have predicted (given the context, including communicative situation, my knowledge of the interlocutor, etc.), that the question concerns lunch *before* the word “lunch” was vocalized. The time pressure of the efficient linguistic communication (not entirely unlike the time pressure of an upcoming Armageddon) forces us to react to an expected result of the process of understanding before the process is completed.

However, the process of *understanding the utterance* does not finish when the process of *uttering the response* begins. This is demonstrated by the fact that when the process of *understanding an utterance* is completed, and the final state of *understanding of this utterance* achieved, we may realize that our prediction has been mistaken. In such a case there are two possible scenarios. Either we have already started to vocalize an inappropriate response and thus we follow it with a repair such as: “I’m sorry, I thought that you were asking about lunch.”; or, given enough time and sufficiently keen reflexes, we might block the process of uttering the response and reconsider our own turn. The necessity to repair one’s own initial reply, or encountering a highly unexpected utterance, will typically result in a considerably longer

gap between subsequent turns (Sacks et al., 1974; Kendrick and Torreira, 2015).

Illocutionary force

There is one more important lesson regarding linguistic understanding that we can learn from the turn-taking studies if we accept LU and extend our characterization of linguistic understanding to account for its processual aspect. Of course, it is a matter of what one's goals are in formulating a theory of linguistic understanding; but it seems that we cannot talk of full understanding of a given utterance unless we mean both understanding of its content and its illocutionary force. Interestingly, this observation finds support in empirical research on turn-taking. According to conversation analysis (Schegloff, 2007; Levinson, 2013), it is impossible to navigate the unfathomable environment of everyday language use and produce subsequent turns with such speed and accuracy, unless we deploy specific strategies and heuristics. One of the most important of these is action ascription.¹⁵ Most discourses are logically organized around so-called "adjacency pairs". For example, if one speaker asks a question, it is highly probable that the turn of the interlocutor will consist of an answer; offers trigger acceptances and

¹⁵Even though the way that conversation analysis characterizes linguistic actions is different from the way philosophy of language typically characterizes speech acts, I take the two enterprises to be complementary. The primary interest of conversation analysis is a statistical classification of countless patterns observable in human linguistic activity (Schegloff, 1992, 2007). Speech act theories, on the other hand, are primarily concerned with constitutive features (either conventional, intentional, functional, expressive or normative) of selected speech acts. Nevertheless, even though the list of speech acts studied in philosophy is shorter than the list of linguistic actions studied in conversation analysis, there can be little doubt that whatever work has already been done in relation to, e.g., asserting, asking, promising, etc., is but the tip of the iceberg of all the speech act types that figure in our everyday linguistic activity. More than anything, it might be a matter of fine-grainedness; maybe speech act theory will come up with a shorter list of more general speech act types that could then be used as scaffolding for the classification of the myriads of subtypes of linguistic actions. As Austin himself already admitted in his foundational work (1962): "I am not suggesting that this is a clearly defined class by any means" (p. 99).

rejections; complaints trigger explanations; etc. In general:

The challenge for participants, then, is to assign at least one major action to a turn they have only heard part of so far. But to do this, they must have parsed what they have heard and understood its grammar well enough to predict both the content and its structure. . . (Levinson, 2013, 103)

The following picture emerges. In spoken linguistic exchanges understanding begins immediately after the hearer’s cognitive system recognizes perceived sounds as linguistic input.¹⁶ The first important task of the hearer’s cognitive system is to obtain sufficient information on the grammatical and lexical features of the utterance to classify it as performing a particular speech act. This action ascription, together with background knowledge, guides further processing of the utterance. For example, after hearing “Would you . . .?”, the comprehension system of an English speaker might initially classify the input as performing the action of asking and use this “assumption” to predict a set of possible grammatical continuations. Obviously, the initial classification is not written in stone. An utterance of “Would you be so kind as to close the window?” has the grammatical structure of a question, but emerging from the mouth of an English speaker would typically perform the action of requesting.

¹⁶As I have suggested above, the process of understanding is involuntary. It is widely recognized that linguistic understanding is not an intentional action (Hunter, 1998) and that it is mostly performed by our cognitive system “sub-personally” and independently of our will (Pettit, 2010). We cannot prevent our understanding of a sentence spoken in a language we understand in the dispositional sense, once we hear it. We cannot decide not to trigger the understanding process. Moreover, even though there is an extensive range of actions that we can perform to improve our understanding (in some cases we can facilitate it by, e.g., repeating an exceptionally complicated sentence aloud or writing it down and parsing it), the occasions on which we actually do so constitute but a tiny portion of our omnipresent linguistic activity. Typically, the understanding process initiates automatically, compulsorily and does not require any personal-level action (see, e.g., Marslen-Wilson and Tyler, 1981; Fodor, 1983).

4 Doing away with the process?

In the previous section, I proposed a definition that outlines the relations between different senses of linguistic understanding (LU) and argued that it is appropriate to characterize linguistic understanding as a process. In this section, I will respond to a possible objection according to which we can do away with the process-talk in the context of understanding.

The following worry may arise. Maybe what I am characterizing as a process of understanding is actually just a sequence of states of understanding? For example, when understanding an utterance “What would you like to have for lunch?” a hearer goes through a succession of mental states representing subsequent fractions of the utterance as it unfolds (e.g., “What...”, “What would...”, “What would you...”, etc.).¹⁷ Thus, we can account for the dynamic aspect of linguistic understanding without characterizing understanding as a process.

I think that this worry is ill-grounded. No doubt that the process I characterize can be alternatively modeled as a succession of states. Plausibly, every cognitive process can be conceptualized as a succession of states leading to its final output. Brian O’Shaughnessy (2000) speaks about it when he contrasts non-experiential processes (such as *forgetting*) with experiential processes (such as *hearing*). While experiential processes are “of necessity constituted of nothing but process-parts” because “[e]verything experiential is of necessity in flux.” (O’Shaughnessy, 2000, 49), non-experiential processes are constituted of both: state-parts and process-parts. O’Shaughnessy explains:

...just as being at position p is one thing and being at p and moving is another, so being in a certain memory-state at time t is one thing and being in that state and forgetting some of that

¹⁷For simplicity I assume here that the utterance is represented in the subsequent states word-by-word. Competing models, e.g., phoneme-by-phoneme, are also possible.

content is another. ... all that is needed for the occurrence of the requisite process is a continuity across time of states which stand to one another in some requisite relation. With this in mind I say, that when the nonexperiential process of forgetting takes processive form, that process is constituted out of memory-states. (O'Shaughnessy, 2000, 46)

Understanding an utterance is a non-experiential process which overlaps with the experiential process of *hearing an utterance*.¹⁸ Both of them are continuous and “exist by unfolding from one time to another” (Crowther, 2009, 5). The fact that we can point at state-parts constituting the process of *understanding* does not make our characterization of understanding as a process redundant or false.

In the existent debate about linguistic understanding, most authors focused on states of understanding of the whole utterance, i.e. the final outputs of the process of *understanding an utterance*. There is a good reason to pay special attention to such states. As highlighted by Longworth: “What we seek in an account of state-understanding is an account of how such states can play a role in ordinary psychology, how occupying them can impact on the rational development of one’s cognitive economy.” (Longworth, 2008b, 51-52). States of understanding interact with various other cognitive states of the subject and thus constitute the link between our linguistic activity and other things we do and believe. For example, I go to the kitchen because I believe you said there is carrot cake in the kitchen; I have a desire to have a piece of carrot cake; I know that the kitchen is close enough for me to reach in just couple of seconds; etc. In principle, there is no reason why theories of states of understanding could not be extended to give an account of state-parts of the process of understanding. However, even equipped with a theory of states of partial understanding leading up to full understanding, we should not resign from describing *understanding an utterance* as a process.

¹⁸For more about the nature of experiential processes see Soteriou (2013, 2018).

Like other non-experiential processes, linguistic understanding “is susceptible of two different analyses, according as different analytical agencies are brought to bear upon it” (O’Shaughnessy, 2000, 45).

5 The “process first” approach to linguistic understanding

My main task in this chapter was to offer a theory of linguistic understanding that would help us account for the dynamic aspect of everyday linguistic interactions; especially the phenomenon of fast turn-taking. I claimed that to do this we should focus on the nature of the process of understanding. However, since philosophers used to focus, up to this point, mostly on the topic of states of understanding, one might have an impression that my discussion in this essay goes orthogonal to the philosophical debate. There are two things to be said about this worry.

Firstly, even though I appreciate and share the interest of other philosophers in the problem of the nature and function of states of understanding, I suggest expanding our attention to other aspects of linguistic understanding. Human linguistic activity is a very complex and multidimensional phenomenon. One of its most striking features is the speed with which we are able to exchange information during everyday conversations. We cannot explain this feature unless we take a closer look at the nature of linguistic processing. Even if typically studied using empirical methods, this process may be of interest to anyone investigating the phenomenon of language use.

Secondly, the existent theories are more or less in agreement regarding a few most important criteria a state has to meet to qualify as a state of understanding. Firstly, it has to be a representational state; representing at least the content (and, possibly, also the illocutionary force) of a given utterance. Secondly, it has to be a personal-level state, i.e., a state consciously accessible to the hearer. Finally, it has to be a state that can interact with other

personal-level representational states of the subject. These are all necessary conditions for something being a state of understanding. But how can we tell which among the competitive theories is the right one?

LU outlines how the different elements of linguistic understanding: the process, the state, and the disposition to understand a given language, are related to each other. States of understanding are generated by the process of understanding; they are the end products towards which the accomplishments *understanding an utterance* unfold over time. If this is the case, investigating the nature of the process might teach us a lesson about the nature of states.

To illustrate it with an example. It is a central point of Longworth’s (2018) theory (in contrast with, e.g., Fricker’s (2003) theory) that states of understanding represent the content of the utterance directly (they are attitudes towards the content of the form “p”), and not *as uttered by a given speaker* (i.e., attitudes towards the content of the form “that S said that p” where S represents a speaker). However, since both these alternatives are compatible with the criteria enumerated above, we might be in a better position to say which one is right if we take a closer look at the mechanisms involved in language processing. For example, the data enumerated by Polka and Nazzi (2018) suggest that:

...although the “who” and “what” sides of the signal are distinct, their processing is functionally integrated, as suggested by research on adults. For example, adults are better at learning to recognize talkers in their native language than in an L2, and dyslexics show impairments in both native language talker recognition and phonological processing... (Polka and Nazzi, 2018, 759)

If the processing of two types of information is in fact integrated, it is not unlikely that it gives rise to a richer type of representation: “that S said

that p”. Longworth suggests that we should not expect, e.g., small children who seem perfectly able to understand simple utterances, to have conceptual resources necessary to represent content involving concepts such as BEING SAID. But here again the discussion is very much open given that the research on speech processing suggests that:

Infants successfully extract information about who is talking early in development (the “who” side of the message), well before they begin to sort out the “what” side of the message, that is, the referential content of words. (Polka and Nazzi, 2018, 758)

I am not claiming that this data alone shows that Longworth is mistaken¹⁹ or constitutes conclusive evidence for any particular stance in the debate. What I want to illustrate is that if we want to choose between different types of representation performing a given task equally well, it might be beneficial to take a look at the way they are actually produced.

6 Concluding remarks

The goal of this chapter was to offer an account of the process of linguistic understanding. Only by focusing on the process sense of understanding, we are able to capture its dynamic and temporal profile. Just like other processes, linguistic understanding unfolds over time.

Crucially, this does not imply that focusing on the state of understanding or the achievement of understanding is inappropriate or misguided. Linguistic understanding is a multidimensional phenomenon and “linguistic understanding”—a capacious notion. As outlined in LU, different aspects of this

¹⁹Longworth’s account of linguistic understanding is very elaborate (Longworth, 2008a, 2008b, 2010, 2018), and I could not hope to address all its subtleties here. In this section, my only goal is to suggest that by paying special attention to the process by which the states of understanding are generated, we might collect additional data worth taking into consideration while investigating the nature of these states. More on these issues in Chapter 2.

phenomenon are closely intertwined. My goal was to illuminate the one aspect relatively absent from the philosophical debate to this date. I hope that this will contribute to our gaining an even better grasp on this complex yet fascinating topic in the future.

Chapter 2

The representational structure of linguistic understanding

Abstract. Building on the strengths of available philosophical accounts, and drawing from empirical research on various aspects of language comprehension, this chapter offers a new account of the representational structure of linguistic understanding. In contrast to existing theories, it postulates that the representational structure of linguistic understanding consists of multiple interdependent representations generated by a dual-stream process. Moreover, the proposed model establishes a middle ground between two competing accounts of the relation between comprehension and acceptance: *Cartesian*, according to which we are free to either accept or reject comprehended contents and; *Spinozan*, according to which we automatically believe whatever we comprehend.

1 Introduction

Linguistic understanding is a central element of our social lives. The fact that the hearer understands what the speaker says is one of the main factors contributing to their successful communication and action coordination. Moreover, linguistic understanding enables the transmission of knowledge through testimony. No doubt linguistic understanding is important, but what is it exactly?

To answer this question, we have to go back to the disambiguation presented in Chapter 1. While talking about linguistic understanding we might have in mind either a *disposition* to understand utterances of sentences in a given language (i.e., knowledge of a given language), a representational *state* that results from exercising this disposition upon encountering a linguistic input (i.e., state of understanding) or, finally, a process that generates this state. This chapter focuses on states of understanding.

There is no consensus in the philosophical literature regarding the best way to characterize the state of understanding. Is it a type of knowledge (Evans, 1982; Heck, 1995); belief (Millikan, 2004, 2005; Balcerak Jackson, 2019); or does it belong to yet another attitude-type (Hunter, 1998; Fricker, 2003; Longworth, 2018)? What exactly is the content of this attitude? In this chapter, I will argue that we should abstain from identifying linguistic understanding with a single mental attitude towards a particular representation. It is much more appropriate to talk about the *representational structure of linguistic understanding*, which consists of multiple interdependent representations. I will demonstrate that some controversies between the available philosophical accounts originate from the exclusive and narrow focus on selected aspects of linguistic understanding. Other controversies, however, turn out to be more substantial. Luckily, we can solve them relying on knowledge from empirical language sciences.

The model I will propose is a result of an inference to the best explanation based on the analysis of available philosophical theories of states of

understanding and data coming from the empirical research on language processing, text comprehension, belief-fixation, and source monitoring. Crucially, this model spells out some major constraints that we ought to take into account while appealing to the nature and function of linguistic understanding in philosophical inquiries into the nature of communication and the epistemology of testimony.

In Section 2, I provide a quick overview of the current philosophical debate about states of linguistic understanding. In Section 3, I discuss a broad spectrum of empirical research that sheds light on the nature of the representational structure of linguistic understanding. In Section 4, I offer a new model of this structure. Section 5 concludes with a discussion of the philosophical relevance of my model.

2 Philosophical accounts of states of understanding

Let us start with a simple story. Rebeka and Hamid are walking along a seashore. Rebeka reminds herself of an interesting fact that the average lifespan of a starfish is 35 years. She wants to share this curiosity with Hamid. In order to do this, she utters the sentence “The average lifespan of a starfish is 35 years.” Hamid is a competent English speaker. Upon hearing Rebeka’s utterance, he recognizes the sounds produced as an utterance of the sentence “The average lifespan of a starfish is 35 years.” He grasps the meaning of her utterance. As a result, he acquires knowledge that Rebeka has said that the average lifespan of a starfish is 35 years. Plausibly, if he did not know it before, he also acquires knowledge that the average lifespan of a starfish is 35 years.

Hopefully, we will be able to identify an element (either explicit or implicit) of this story as the target of our investigation, i.e., the state of linguistic understanding. In order to presuppose as little as possible, I introduced

the scenario using vague (e.g., *recognizing* sounds as linguistic utterances; *acquiring* various forms of knowledge) or even metaphorical (e.g., *grasping* meaning) terms. Since all of these terms call for elucidation, and since it cannot be offered in a theory-neutral way, I will use this story to review the most influential among the philosophical theories of the state of understanding. Borrowing a helpful distinction from Guy Longworth (2018), we can divide them into two groups: *indirect* and *direct*.¹ The distinction is based on the way in which a state of understanding postulated by a given theory represents the content of the linguistic input. According to indirect views, the content (p) of an utterance is represented *as being said/asserted by a given speaker*.² When Rebeka says that p , Hamid’s understanding has a content: $r = \text{that Rebeka said/asserted that } p$. According to direct views, on the other hand, the content of a state of understanding is just the content of the utterance, i.e., p .

There are three most prominent indirect views: *the knowledge view*, *the linguistic seeming view*, and *the indirect belief view*. According to *the knowledge view*, Hamid’s understanding of Rebeka’s utterance is identical with his propositional knowledge *that Rebeka said that the average lifespan of a starfish is 35 years* (Evans, 1982; Heck, 1995). The knowledge state represents the content of Rebeka’s utterance indirectly, i.e., as said by Rebeka. Obviously, we could not demand from a comprehender to have a direct knowledge of whatever was the content of an understood utterance; one can understand a false utterance, but one cannot know what is not true.

The second indirect view is *the linguistic seeming view* formulated by Elisabeth Fricker (2003). Fricker characterizes linguistic understanding as a type of seeming (she uses the term *quasi-perception*) of the content and force of an utterance.

¹Longworth uses different terminology; he speaks about *first-* and *second-order states of understanding*.

²Given that indirect representations have direct representation embedded in it, I will also refer to indirect representations as *meta-representations* (cf. Sperber, 2000).

In fact she [the hearer] may not, as it were, bother to form a belief about what was said at all. The most immediate personal-level psychological effect of her auditing of the utterance is that she enjoys a representation of a distinctive kind special to language understanding: a conscious representation of the content and force of the utterance. (Fricker, 2003, 325)

When Hamid hears Rebeka’s utterance during their walk, it seems to him *as if Rebeka asserted that the average lifespan of a starfish is 35 years*.³ Again, as in the case of the knowledge view, the content of the seeming is not the same as the content of the utterance. It does not seem to Hamid that the average lifespan of the starfish is so and so, only, that Rebeka said that it is so and so.⁴

The third of the indirect views is the *indirect belief view* (cf. Balcerak Jackson, 2019). According to this view, Hamid’s understanding of Rebeka’s utterance is identical with his belief that Rebeka said that the average lifespan of a starfish is 35 years. Again, just like in other indirect views, the state of understanding meta-represents the content of the target utterance as coming from a given source.

On the other side of the barricade, we find two prominent views according to which states of understanding are direct (i.e., represent exclusively the content of the target utterance): *the content-entertaining view* and *the direct belief view*. According to Guy Longworth’s (2018) *content-entertaining view*,⁵

³On Fricker’s account, seemings are *prima facie* justificatory for beliefs about the force and content of the speaker’s utterance. For a discussion see Balcerak Jackson (2019) and Chapter 3 of this thesis.

⁴A related view has been offered by David Hunter (1998). Hunter identifies linguistic understanding with an experiential state “of immediate awareness... of the text’s or speaker’s meaning” (p. 577).

⁵What I describe here is Longworth’s theory of linguistic understanding as presented in his (2018). His full theory of language comprehension is fairly nuanced (and in many respects similar to my positive proposal). Crucially, he distinguishes *comprehension* from *understanding* and treats comprehension as a super-faculty constituted by two sub-faculties: understanding, and speech perception (Longworth, 2008a, 363). More on this below.

to understand an utterance is to entertain the very proposition expressed by this utterance (p), and not some other proposition *about* p , e.g., $r = \textit{that Rebeke said that } p$. On this account, Hamid’s understanding of Rebeke’s utterance equals his *engaging* with or, *entertaining* that the average lifespan of a starfish is 35 years (as a result of perceiving Rebeke’s utterance). Importantly, according to Longworth, entertaining does not entail acceptance. One can entertain given content without believing it.

This last assumption is rejected by the second direct view, i.e., the *direct belief view* (Millikan, 2004, 2005; cf. Mandelbaum, 2014; Mandelbaum and Quilty-Dunn, 2015; Recanati, 2002). According to this view:

We do not first understand what is said and then evaluate whether to believe it. Rather, we first believe what is said and then, if we are not under too much cognitive stress, we may think it over critically and reject it. (Millikan, 2004, 121)

Normally, hearing that p is believing that p —possibly not in the purely statistical sense of “normally” (although I would argue that that is so too), but believing that p is the default. It is what happens when nothing intrudes. (Millikan, 2004, 129)

According to this account, for Hamid to understand Rebeke’s utterance is to form a belief that the average lifespan of a starfish is 35 years.

To sum up, we can distinguish five views about the nature of states of understanding. Three indirect: the knowledge view, the seeming view, and the indirect belief view; and two direct: the content-entertaining view, and the direct belief view. Despite obvious differences, all these views have a lot in common. In the remainder of this section I would like to indicate the five most important features of the state of understanding accepted by all these accounts.

Firstly, the theories enumerated above agree that understanding is a representational state.

Representation: the state of understanding represents the content (and, plausibly, the illocutionary force) of an utterance.

A state of understanding represents the content or *what is said* in an utterance either directly or indirectly. Whether it also represents the illocutionary force of this utterance is more debatable. The force is represented in states of understanding according to indirect views. In general, a hearer who understands an utterance has knowledge/belief/seeming that the Speaker asserted/asked/ordered/etc. that p. As it stands, this view is controversial because it presupposes that comprehenders have conceptual resources necessary to represent the force of speakers' utterances. A more compelling story might be available; however, it merits a separate discussion. In this chapter, I will focus on the (complicated enough) topic of understanding the content of assertoric speech.

Secondly, competing theories converge on the claim that states of understanding are consciously accessible to the receiver.

Accessibility: the state of understanding is, in principle, accessible to comprehender's consciousness (i.e., it is a personal-level state).⁶

I will follow Rowlands (2010) in distinguishing two main types of cognitive processes. *Personal-level* cognitive processes are ones which make information (unavailable prior to this processing) available to the subject. *Sub-personal* cognitive processes are ones that make information (unavailable prior to this processing) available to subsequent processing operations but *not* to the subject. The process of language comprehension is a paradigm examples of personal-level process, i.e., a process that makes previously unavailable information about the content of a linguistic input available to the comprehender.

⁶This and the next feature of the state of understanding are pointed out by Longworth (2008b, 51).

Thirdly, understanding can (and often does) interact with other personal-level states:

Interaction: the state of understanding can interact with other personal-level states, such as beliefs, knowledge, and desires.

The most important role of states of understanding is to link our linguistic activity with other things we think and do. If you tell me “There is carrot cake in the kitchen” and I jump up off my chair and run out of the room, I plausibly do it *because* I understood that you have told me that there is carrot cake in the kitchen. My understanding of your utterance *interacts* with (i) my desire to have some carrot cake; (ii) my belief that the kitchen is a place I can reach in just a few seconds, etc.

Additionally, all available accounts agree that an important role of linguistic understanding is that it enables us to acquire two types of knowledge:

Knowledge-enabling: the state of understanding enables comprehenders to acquire (i) knowledge about what other people say and, (ii) testimonial knowledge about the world.⁷

Importantly, at least in the direct views, it is assumed that (i) is not a prerequisite for (ii). However, available theories rarely elaborate on the exact relationship between these two types of knowledge. It is one of the advantages of the model I propose in Section 4, that it allows us to tackle this difficult problem.

The last feature of states of understanding is sparsely discussed in the debate but relatively uncontroversial. It connects three senses of understanding mentioned in the Introduction.

⁷According to the knowledge view linguistic understanding not only enables but is identical with (i).

Origin: the *state* of understanding originates from the *process* of understanding based on the *disposition* to understand utterances of sentences of a given language (i.e., linguistic competence).⁸

A state of, e.g., entertaining the content expressed through a given utterance does not qualify as a state of understanding this utterance, unless it was generated by an appropriate type of processing of this utterance. I can entertain the content of an utterance u without even being aware that u was ever uttered and thus without a linguistic understanding of u . This is the case, for example, if I have understood another utterance w (possibly in a different language), which expresses the same content as u .⁹

These five general features of states of linguistic understanding demarcate the subject area of our investigations. By keeping them in mind we will avoid the risk of changing the topic of our discussion while engaging with the empirical literature on language comprehension in the rest of the chapter.

To sum up. The intuitions driving the accounts enumerated in this section seem to pull us in different directions. Do we automatically believe everything we comprehend? Do we typically follow just *what* is said or, also, *who* said that? My goal in the rest of this chapter is to offer an account that allows us to answer these questions and reconcile the intuitions voiced by different contributors to the debate. In the next section, I take a closer look at the empirical research on language comprehension to find out what it teaches us about the types of representations employed during linguistic understanding.

⁸Cf. my discussion in Chapter 1.

⁹More on this topic in Chapter 5.

3 Psychology of language comprehension

3.1 From propositional representations to situation models

Looking for psychological counterparts of mental attitudes postulated by philosophical theories of understanding, one undoubtedly comes across the classical Construction-Integration model of language comprehension (Kintsch and van Dijk, 1978; van Dijk and Kintsch, 1983; Kintsch, 1988).¹⁰ The model distinguishes three levels of representations, two of which are particularly relevant to our discussion: *propositional textbase representation*, and *situation model*.¹¹ The construction of a propositional representation is postulated as an initial step of text processing.

The textbase level is represented in terms of propositions. One important assumption of the model is that the fundamental unit of processing is the proposition, which consists of a predicate and argument(s). The proposition generally represents one complete idea. (McNamara and Magliano, 2009, 309)

At the first stage of processing, comprehension generates propositional representations of bits of discourse.¹² These representations are encoded in

¹⁰In this section, I will be focusing on models of text comprehension. It is commonly believed that assumptions about processing of written text extend to comprehending speech. Although, due to relative difficulty of empirical testing of speech comprehension in comparison with text comprehension, the body of empirical research on spoken discourse comprehension is still small (but see, e.g., Piest et al., 2018).

¹¹The third level of representation is the *surface structure* representation, i.e., the result of “decoding of phonetic and graphic strings, the identification of phonemes/letters, and the construction of morphemes” (van Dijk and Kintsch 1983, 13).

¹²It is worth noticing that the characterization of propositions used by van Dijk and Kintsch is closely related to a fairly standard philosophical characterization: “a proposition is an abstract, theoretical construct, which is used to identify the meaning, or what is expressed by a sentence under specific contextual restrictions (speaker, time, place), and which is related to truth values” (van Dijk and Kintsch, 1983, 111).

the episodic memory and become, at least to some extent, available to the hearer or reader (cf. **Accessibility**). The formation of propositional representations of linguistic input is an important element of discourse processing, however, the authors of the model make the following reservation:

One must, however, guard against the view that they [propositional representations] are allpurpose representations, and, in particular, provide “the” representation of meaning. (van Dijk and Kintsch, 1983, 38)

They highlight that:

A major feature of our model is the assumption that discourse understanding involves not only the [propositional] representation of a textbase in episodic memory, but, *at the same time*, the activation, updating, and other uses of a so-called situation model in episodic memory: this is the cognitive representation of the events, actions, persons, and in general the situation, a text is about. (van Dijk and Kintsch, 1983, 11; *emphasis mine*)

Even though most of the contemporary theories agree that propositional representations play some role in text comprehension, it is the notion of *situation model* that made a career in the literature on language processing in the last thirty five years (Johnson-Laird, 1983; van Dijk and Kintsch, 1983; cf. McNamara and Magliano, 2009). So, what are situation models?

Situation model is a complex representation of a state of affairs described by a given piece of discourse and constantly updated upon reading subsequent clauses and sentences (cf. Zwaan et al., 1995; Zwaan and Radvansky, 1998; Wyer and Radvansky, 1999; Zwaan, 2016; Richter et al., 2009; Richter and Singer, 2018). Construction of a situation model requires integration and elaboration of information presented across sentences and, thus, engages readers’ background knowledge. Importantly, the formation of a situation

model is often taken to be a measure of the depth of comprehension (Long and Freed, 2018). Researchers focus on five dimensions integrating the information represented in situation models: space, time, causality (i.e., causal relations between events), intentionality (i.e., intentions guiding agents' actions), and entities (either agents or objects) (cf. Zwaan and Radvansky, 1998; Wyer and Radvansky, 1999).

As complex representations, situation models are itself constituted of simpler representations. It is a common assumption in the empirical literature that comprehension delivers representations of multiple formats, all of which contribute to a situation model.

... comprehension involve multiple modalities and modes of representation, including verbal, symbolic, and iconic modalities. The creation of images while understanding discourse is assumed to be fundamental process of comprehension... Certainly, all comprehension models assume that the reader is creating iconic images while reading... The alternative notion is simply ridiculous. (McNamara and Magliano, 2009, 348)

Although the topic of the plurality of representational formats involved in language comprehension is intriguing,¹³ I have to leave it for another occasion. Here, when talking about situation models, I will be referring exclusively to its doxastic component, i.e., beliefs constituting situation models.

Given that language comprehension produces both the propositional representations of pieces of linguistic input and situation models integrating information expressed in these pieces, we have to ask: "How are the two related?" According to the psychology of text comprehension, propositional representations of currently read sentences or clauses are *updated* or *incorporated* into the situation model. The process is cyclical and occurs when-

¹³See, for example, Glenberg and Kaschak (2003) and Shepherd (2018) for the discussion of a possibility that language comprehension produces, among others, motor format representations appropriate to guide physical actions.

ever a representation of a short sentence or phrase is generated (Kintsch, 1988). During processing of a compound sentence (even outside the context of a longer text or discourse), a reader’s cognitive system generates multiple propositional representations (one for each clause or phrase), and instantly integrates each one of them into the situation model of the state of affairs described in the sentence. It is, thus, difficult to point out *the* single state of understanding of utterances of even such apparently uncomplicated sentences as “I really need to go to work, but I am too sick to drive.” According to empirical models of comprehension, while comprehending an utterance of this sentence, a reader’s cognitive system generates two propositional representations (one for each clause) which are subsequently integrated into a situation model representing the described state of affairs. Moreover, if the next sentence says: “I have to call a taxi or ask Mark to give me a lift,” propositional representations generated as a result of processing this input get integrated with the same situation model.¹⁴ Crucially, when asked about the content of a comprehended bit of discourse, readers retrieve information primarily from the situation model. It is assumed that memory traces of propositional textbase representations are consulted only when answering questions regarding which information was *explicitly* stated in the text, and are quite limited (Sparks and Rapp, 2011; Long and Freed, 2018). Finally, it is a crucial assumption of the leading psychological models of comprehension, that the information from propositional representations is updated to the situation model *automatically*. Does it mean that we automatically believe whatever we comprehend? In the next section, I take a closer look at this exact question.

¹⁴Obviously, it is possible to be building two (or more) distinct situation models at once. This happens, for example, if one reads a book about the World War II while simultaneously listening (or half-listening) to her friend describing his day at work.

3.2 Comprehension, acceptance, and validation

Motivating her theory of understanding as a direct belief, Millikan (2004, 121, 2005, 117) appeals to research of a psychologist Daniel Gilbert.¹⁵ In the early ‘90s Gilbert and his colleagues published several studies designed to test two competitive models of acquisition of beliefs through linguistic comprehension (Gilbert et al., 1990; Gilbert, 1991; Gilbert et al., 1993). According to the so-called *Cartesian* model, people do not believe everything they comprehend. Understanding is separated from acceptance; when we comprehend an utterance, we first entertain a proposition it expresses and then, in a subsequent step, either accept or reject it. This view is in line with Longworth’s (2018) content-entertaining view. Discussing an example of understanding “Smoking is dangerous,” Longworth says:

Although understanding such an assertion involves engaging the content that smoking is dangerous at first order [directly], it does not—or need not—involve accepting that content. (Longworth, 2018, 824)

According to the second model discussed by Gilbert, the so-called *Spinozan* model, understanding is believing. Acceptance is a default position towards comprehended content, while rejection is an effortful activity, which requires time and cognitive resources, and happens only after a belief has been already acquired (cf. Mandelbaum, 2014; Recanati, 2002). Loosely speaking, we cannot prevent the content we comprehend from getting into our belief-box. All we can do is try and get rid of it once it is already there.¹⁶

¹⁵Cf. Kissine and Klein (2013), Mandelbaum (2014), Mandelbaum and Quilty-Dunn (2015).

¹⁶Gilbert (1991) mentions also a hybrid *Cartozan* model according to which we are biased towards acceptance (acceptance and rejection are not on par), but *mere* comprehension and acceptance are separate states. However, Gilbert (1991) argues that there is no clear empirical evidence for the existence of *mere* comprehension as separate from acceptance. Probably this is why Cartozan model is virtually absent in subsequent literature on this topic.

Gilbert argued that the Spinozan model is correct. One of the first studies designed to establish it was “The Hopi Language” experiment (Gilbert et al., 1990, Study 1). Participants of this experiment read a series of statements of the form *An X is a Y*, with an English noun in place of *Y* and what they were supposed to believe is a Hopi word (in fact a nonsense string of letters) in place of *X* (e.g., *A tarka is a wolf*). Some of the statements were followed by a display of the word *true* (indicating that the previous sentence was true), others with *false*, yet others with no display of either *true* or *false*. Most importantly, during the presentation of some statements, participants heard a tone. Earlier, they had been instructed to press a button as quickly as possible each time they heard it. This “interruption task” was intended to increase the cognitive load and make the subjects’ processing of the information on the screen more challenging.

The initial learning phase was followed by a testing phase where participants were asked about the meanings of words, which they had learnt through the statements presented in the learning phase (*Is X a Y?*). According to Gilbert and his colleagues, if it would turn out to be the case that the additional cognitive load makes it difficult to tag statements as false but does not influence tagging statements as true, the Spinozan model would be vindicated; accepting comprehended propositions is automatic but rejecting them is costly and requires effort. The results revealed the expected pattern and supported the Spinozan model.

It is, however, unclear whether this particular experiment actually tested comprehension-based belief-fixation or just memory of the learned information (cf. Kissine and Klein, 2013). To make sure that the Spinozan model is correct, Gilbert and his colleagues conducted another set of experiments (Gilbert et al., 1993). In one of them, participants read crime reports consisting of multiple statements, some of which were false (as indicated by use of the red font) while others were true (displayed in black font). As in the Hopi Language experiment, researchers used an additional interrup-

tion task, to selectively increase cognitive load. Crucially, after the learning phase, participants were asked not only memory retrieval questions (about whether a given statement appeared as true, as false, or did not appear at all). They were also asked, for example, to recommend a prison term for the perpetrators based on the reports just read. The results showed that: “Interrupted subjects recommended that perpetrators serve nearly twice as much time when the false statements contained in the police reports exacerbated (rather than extenuated) the severity of the crimes.” (Gilbert et al., 1993, 225). Again, it seems that—as predicted by the Spinozan model—participants initially accepted all information as true. Those of them, who were simultaneously distracted by an interruption task, did not manage to remove the propositions explicitly indicated as false from their belief box. As a result, they kept assuming this information to be true while recommending a prison term.

As far as the results of Gilbert’s experiments go, Millikan’s direct belief view of understanding looks convincing. However, the last 25 years of research on the relation between comprehension and acceptance push us towards a more nuanced picture. Hasson et al. (2005) conducted two experiments designed to test the Spinozan model. In the first experiment, they demonstrated that *highly informative* false sentences were not represented as true even if participants operated under cognitive load. Gilbert and his colleagues’ (1990) results were replicated only in relation to *uninformative* false sentences. In the second experiment, Hasson et al. demonstrated that in a lexical decision task participants were considerably faster to associate an adjective (e.g., *optimist*) with someone who has been previously characterized with a matching statement (e.g., *this person thinks that things turn out for the best*) if the statement was said to be true of a person in comparison to it being said to be false or neither true nor false. This suggests that, in contrast with Gilbert’s assumption, sentences marked as false and sentences whose veracity is unknown might not be automatically encoded as true: some

degree of belief suspension is possible (*see also* Street and Richardson, 2015). Finally, Pantazi et al. (2018) suggested that the tendency to believe comprehended information (both true and false) might be independent of cognitive load. This result also undermines the Spinozan model whose central assumption was that subjects are unable to filter out false information *only if* their cognitive resources are limited, i.e., under cognitive load.

All these results suggest that a third way between Spinozan and Cartesian view of language comprehension is needed. To find it, we should focus on the research on the mechanism of so-called *validation* (cf. Wyer and Radvansky, 1999; Singer, 2006, 2013, 2019; Schroeder et al., 2008; Richter et al., 2009; Isberner and Richter, 2014; Richter, 2015; Kendeou, 2014; O’Brien and Cook, 2016). Validation is a process of monitoring incoming information both for internal consistency and consistency with comprehender’s knowledge (Richter et al., 2009).¹⁷ If a comprehender has relevant background information, either active in working memory or easily accessible in long-term memory, this information will be used to validate the content of a linguistic input independently of cognitive load.¹⁸ Minimally, validation is assumed to detect “violations of factual world knowledge (e.g., *Soft soap is edible*), implausibility (e.g., *Frank has a broken leg. He calls the plumber*), inconsistencies with antecedent text (e.g., *Mary is a vegetarian... She orders a cheeseburger*), and semantic anomalies (e.g., *Dutch trains are sour*).” (Isberner and Richter, 2014, 246).¹⁹

¹⁷This mechanism is similar to what Sperber et al. (2010) call *vigilance towards the content*.

¹⁸This is compatible with results of Gilbert’s Hopi Language experiment where participants lacked background information against which the linguistic input could be validated.

¹⁹I focus on comprehending discourses that are *not* established or recognized as fictional. In the case of comprehending fiction, validation is *recalibrated*: “. . . specific discourse contexts, most notably stories that create a fictional story world, seem to modulate validation to some degree.” (Richter and Singer, 2018, 184). Situation models constructed while comprehending fiction contain beliefs about the world of the fiction. Therefore, if it is sufficiently obvious that we speak about a fictional world, e.g., *the Incredible Hulk has thrown a lorry*, can successfully pass validation and get integrated with the situation model.

Validation is *routine* and *non-strategic*, i.e., independent of subject-specific processing goals.²⁰ For example, Richter et al. (2009) demonstrated that participants asked to monitor the orthographical correctness of a statement, produced affirmative orthographical judgment (that the target word in the statement was spelled correctly) slower and with lower accuracy when the target statement was false.²¹

Crucially, validation serves as a precondition on updating a situation model with the information captured in propositional textbase representations (Wyer and Radvansky, 1999). According to Schroeder et al. (2008), the construction of a situation model is guided by two main directives: *accuracy* (to represent the state of affairs described in the discourse as accurately as possible) and *stability* (to represent the state of affairs described in the discourse in a stable and consistent way).

How do comprehenders manage to achieve both accurate and stable representations? We suggest that they carry out epistemic validation processes that monitor whether incoming information is consistent with other ideas provided in the text, with the current state of the situation model, and with general world knowledge. We assume that these validation processes are routinely carried out when situation models are updated and that they are a major determinant of whether a particular piece of information is integrated into the situation model, with the potential consequence of altering a comprehender’s world view. (Schroeder et al., 2008, 238)

The fact that the formation of situation models is guided not only by accuracy but also stability has some troublesome consequences from the epistemic

²⁰Comprehenders *can* engage in validation strategically (cf. Singer, 2019) but not all kinds of validation are strategic and intentional. Here, I focus on the routine validation (sometimes called *epistemic monitoring* (cf. Schroeder et al., 2008)).

²¹See Isberner and Richter (2013, Experiment 2) for similar effect obtained using non-linguistic, color judgment task.

point of view. According to Schroeder et al. (2008), once information passes the gatekeeper of validation and gets integrated with the situation model, it becomes a background for further validation of new information. By the same token, it becomes very difficult to remove, even if it turns out to be false. Validation is by no means a perfect counter-deceptive tool. It has considerable limitations: it employs only *available* and *activated* background information (which may in itself be false); it often fails to filter out sufficiently plausible false information; it is based only on quick and incomplete analysis; etc. (Isberner and Richter, 2014; Richter, 2015; Marsh et al., 2016).

The following picture emerges. Linguistic comprehension generates (at least) two types of representation: propositional representations and situation models. When we understand an utterance of a compound sentence (not to mention a bigger piece of discourse), propositional representations of particular phrases and clauses get cyclically integrated with the situation model. The update is conditional on their passing by the gatekeeper of validation.²² Validation is routine and non-strategic, i.e., it does not depend on the subject's conscious decisions, their goals, or evaluative mindset. Therefore, neither the content-entertaining view, which follows the Cartesian model nor the direct belief view, inspired by the Spinozan model, got it quite right. We do not automatically believe *everything* we are told, but we automatically believe everything that is not filtered out by validation.²³

²²It is still to some extent an open question what happens to information that does not pass the gatekeeper of validation. One hypothesis is that if the falsehood is informative and the subject knows that it is false, they represent it in terms of what its falsity implies. If I learn that *Lin does not have a sister* is false, I might update my situation model with the information *Lin has a sister*. However, if I learn that *Lin is a pilot* is false, I probably update my situation model with the information *Lin is a pilot* with a *falsity tag*. Apparently, this second way of updating makes me much more prone to the error of misremembering the information as true (which helps explain the results of Gilbert's Hopi language experiment) (Hasson et al., 2005).

²³For a similar assessment of the relation between Spinozan and Cartesian models see (Kissine and Klein, 2013).

3.3 Vigilance towards the source

Thus far I did not say anything about representing the content of a linguistic input *as said by a particular speaker* or *as coming from a particular source*. Does it mean that indirect views of linguistic understanding are totally off the mark? I do not think so. Representing a given content as coming from a particular source is one of the crucial skills of comprehenders.

The literature on the phenomenon of so-called *source monitoring* or *vigilance towards the source* is vast, and I will not attempt to review it in this chapter.²⁴ Different aspects of this phenomenon were to this day investigated in hundreds of empirical studies and incorporated in dozens of theoretical models (for reviews see, e.g., Johnson et al., 1993; Sperber et al., 2010; Mercier, 2017). The core assumption of this research program is that language speakers routinely track the identity of their informational sources, which, in consequence, allows them to monitor the *competence* and *benevolence* of these sources. Some of the hottest questions in the debate concern: the exact mechanisms employed by vigilance towards the source, the effectiveness of our monitoring mechanisms, and the development of vigilance during human ontogenesis (i.e., at what age children become vigilant language users).

Research on source monitoring offers a straightforward lesson about language comprehension: we routinely track which information comes from whom. It is a separate (and somehow scholastic) question whether source monitoring *accompanies* linguistic understanding or *belongs* to it as a part and component. This question turns out to be a bone of contention between the direct and indirect views on the nature of states of understanding. If, as is the case according to indirect views, the state of understanding represents that *a given speaker said that p*, then monitoring turns out to be a component of understanding. If, however, as is the case according to direct views, the state of understanding represents simply the content of an utterance (*p*),

²⁴More on this topic in Chapter 4.

source monitoring *accompanies* linguistic understanding.²⁵

Independently which way we will ultimately go, given what we know about source monitoring in language comprehension, we can assume that an indirect or meta-representation of an utterance as coming from a given source is yet another element of the representational structure generated and employed during linguistic understanding.²⁶ We can think about this representation as having two slots: SOURCE(CONTENT). The slot for SOURCE is filled with a more or less fine-grained representation of the source available to the comprehender, e.g., *someone in the crowd; my mum; the author of this article; the girl with a weird accent*, etc. The slot for CONTENT is filled with a representation of the utterance content similar to the propositional representation.²⁷

²⁵The latter option is endorsed by Longworth (2018). He argues that one does not accept the content one entertains if “. . . one has, or if one takes oneself to have, sufficiently weighty reasons for doubting the competence or sincerity of the assertor.” (Longworth, 2018, 824). However, on his account, source monitoring accompanies, instead of being a part of, understanding. This is a consequence of Longworth’s (2008a) distinction into understanding and comprehension. According to this distinction, comprehension is a super-faculty consisting of an intellectual faculty of linguistic understanding, and perception. The information about the source is not represented in the state of understanding but it is delivered by perception: “The outputs of the super-faculty are cognitions to the effect that that particular utterance gave expression to a particular entertained content. . . .” (Longworth, 2008a, 363).

I think that Longworth’s theory of comprehension is on the right track, but he underestimates to what degree his *understanding* (i.e., the intellectual component of comprehension) is integrated with perception (see discussion in Drożdżowicz, 2019). As indicated by Falandays et al. (2018): “. . . speech perception is readily influenced by lexical and semantic context. . . [t]he current state of the literature now definitively points to a highly parallel, interactive architecture of speech perception.” (p. 2). For example, it was demonstrated that ambiguous acoustic stimuli are perceived differently depending on the meaning of their antecedents, e.g., after comprehenders hear “the man put on his. . .” the ambiguous acoustic token is much more often perceived as “coat” than “goat” (Borsky et al., 1998). Therefore, I do not subscribe to Longworth’s distinction into comprehension and understanding and I will keep using these two terms interchangeably. If someone wants to preserve the distinction, they should think about my theory of understanding/comprehension as corresponding with Longworth’s theory of comprehension.

²⁶See Sperber (1997, 2000) for a discussion of the role of meta-representations in language comprehension.

²⁷As mentioned in Section 2, many indirect views suggest that these representations

4 The representational structure of linguistic understanding

It is time to take stock. In Section 2, I enumerated five theories of states of understanding. According to three of them, states of understanding represent the content of the utterance indirectly (i.e., meta-represent the content) as coming from a given source. According to the remaining two, while understanding an utterance, we engage with its content directly. The empirical data reviewed in Section 3 does not favor either of these approaches. Apparently, to explain what happens when we understand language, we have to postulate at least three types of representations: (i) direct propositional representations; (ii) indirect meta-representations of the content as coming from a given source; and (iii) situation models, i.e., complex representations consisting of beliefs about (and other, non-discursive representations of) the states of affairs described in a given piece of discourse. But how are all these representations related to each other? Is it possible to fit them all into a single, coherent picture? I believe it is.

I would like to offer the following new model of the representational structure of linguistic understanding (fig 1.).²⁸ According to this model, the process of understanding runs in two streams. The first, *direct* stream is faster. It involves the production of the propositional textbase representation which, if it passes the filter of validation, gets automatically integrated into the situation model and shapes our beliefs about the world.²⁹ This stream is responsible for the aspect of linguistic understanding highlighted by the direct

have a further slot for the illocutionary force: SOURCE(FORCE(CONTENT)). I find this view worth exploring but cannot do it in the present discussion.

²⁸As mentioned in the introduction, this model is a result of an inference to the best explanation based on empirical research discussed in Section 3, and intuitions voiced by various philosophical accounts (Section 2).

²⁹As I mentioned above, it shapes our beliefs about the world when the comprehended linguistic input is *not* recognized as fictional. In the case of comprehending fiction, it shapes our beliefs about the world of the fiction.

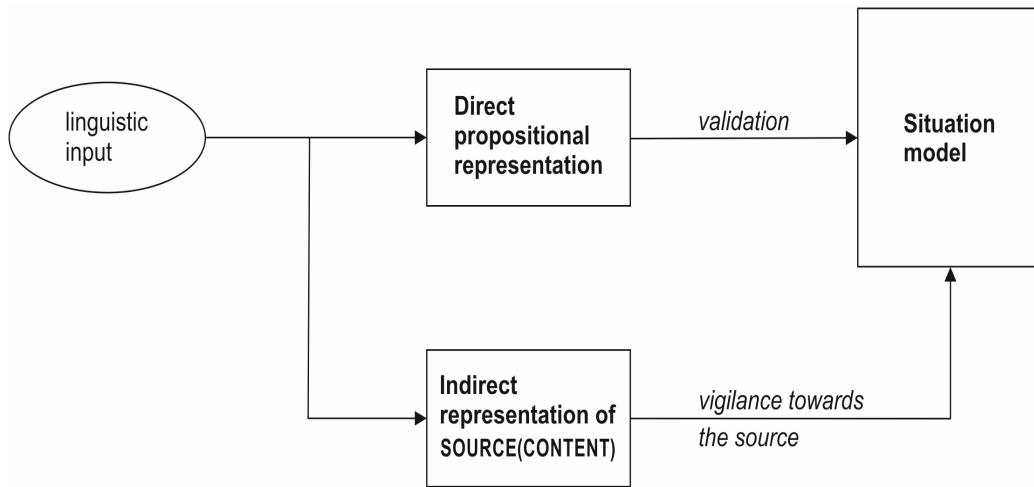


fig. 1 The representational structure of linguistic understanding

views: upon understanding a linguistic input, we immediately and directly engage with its content. When you tell me that a given car is brand new, it makes me initially think about the car (that it is brand new) and not about you (that you have said that the car is brand new).

The second, *indirect stream* is slower. It generates a meta-representation of the content of an utterance as uttered by a given speaker or as coming from a given source. This representation is filtered by vigilance towards the source, which serves as an additional gatekeeper of the situation model.³⁰ If

³⁰While in the case of validation it is clear that at least some validation is routine and non-strategic, and, thus, can be conceived as a subprocess of the process of comprehension, it is more questionable when it comes to vigilance towards the source. Plausibly, we could localize different vigilance mechanisms on a scale from more routine and automatic (e.g., vigilance based on a foreign accent; vigilance based on an inappropriate use of words) to more strategic (e.g., vigilance towards the source who has proved to be unreliable in the past; vigilance towards the source about whose potential deceitfulness we have been warned) (cf. the distinction into type 1 and type 2 vigilance in Michaelian (2013)). Some people might be inclined to characterize the *automatic* vigilance as a part of comprehension

the filter detects that the information comes from an unreliable source, the subject can *attempt* an update of the situation model and belief revision.³¹ If, for example, I know that you are very desperate to sell me the car, I may end up revising my belief that it is brand new.³²

Importantly, the hypothesis about asynchrony of the two streams is, at least initially, vindicated by empirical research. Relying on the results of Sparks and Rapp (2011) and Nadarevic and Erdfelder (2013), Weil et al. (2019) suggest that "...readers consider the credibility of a source only after they have comprehended information and evaluated its consistency with the active memory contents. Accordingly, source credibility might not influence the initial encoding of the information, but rather, encoding might be modified after validation is completed" (p. 231).

It is worth highlighting that the current model is a dual-stream model and not a default-interventionist dual-process model (cf. Gawronski et al., 2014). Comprehension routinely runs through both the direct and indirect stream. The only exception might be very early childhood; it is possible that in human ontogeny the direct stream develops first,³³ and that there is a period in which children are already able to understand language but not yet able to monitor the source of comprehension-based information. Still, the claim is by no means obvious as we have evidence of source monitoring and source identification already in infants (Polka and Nazzi, 2018), and we know that children as young as 3-years-old display some selective trust based on

process and the *strategic* vigilance as a mechanism accompanying comprehension. Others might be inclined to treat all vigilance mechanisms as accompanying comprehension. Nothing important for my discussion hinges on resolving this matter in one or the other way.

³¹Often, after the indirect SOURCE(CONTENT) representation passes through the filter of vigilance towards the source, the source of information gets forgotten (Begg et al., 1992; cf. Michaelian, 2010).

³²More on this topic in Chapter 4.

³³Although I do not want to engage in evolutionary speculations, I think that it is also quite likely that the direct stream appeared first in human *phylogeny* as evolving from older, perceptual mechanisms (cf. Kissine and Klein, 2013).

source identity (cf. Hermes et al., 2018). The question remains opened until more empirical evidence is collected.

5 Conclusion: consequences for the philosophical debate

The model I offered, while compatible with the results of empirical research, reveals that both direct and indirect views point out important aspects of the phenomenon of linguistic understanding. Direct views are right that upon hearing or reading an utterance we directly engage with its content. I argued that we do it through situation models built primarily on the basis of propositional representations. Simultaneously, indirect views are right that we monitor the source of a given linguistic input. If the representation of content *as coming from a given source* does not pass the filter of vigilance towards the source, we make an attempt to update the situation model by removing the information coming from an unreliable source. One thing does not exclude the other—comprehension flows in two streams.

Secondly, the model I offer helps to spell out the complicated relation between comprehension and acceptance. Neither the content-entertaining nor the direct belief view captures this relation well. As demonstrated in the research on validation, even though we do not automatically accept *everything* we hear or read (e.g., validation filters out inputs that are in contradiction with relevant background information active in our working memory or easily accessible in long term memory), we cannot freely choose what to believe and what to reject.

Thirdly, this model has important consequences for the debate about the epistemic role of linguistic understanding. As I said in the beginning (Section 2), one thing that makes linguistic understanding philosophically interesting is that it enables the acquisition of two types of knowledge: knowledge about what other people say; and knowledge about the world based on what other

people say (testimonial knowledge). In most theories it remains unclear how exactly these two types of knowledge relate with each other. The model I offer allows us to spell out this relationship.

The knowledge about what other people say originates from the indirect stream of comprehension, which involves the representation of content as coming from a given source. How it happens can be explained in different ways, and I discuss it in detail in Chapter 3. In a nutshell, we can either assume that our attitude towards the indirect representation of SOURCE(CONTENT) is a belief (cf. Balcerak Jackson, 2019) or some other attitude, e.g., a seeming (Fricker, 2003). If it is a belief, we have to explain how it is justified and whether it amounts to knowledge about what other people say. If it is a seeming, we could argue—as Fricker (2003) does—that it provides *prima facie* justification for beliefs, which (plausibly under further conditions) amount to knowledge about what other people say.

When it comes to the knowledge about the world based on other people’s say-so (testimonial knowledge), the whole dual-stream process of comprehension (with all the representations involved) contributes to its acquisition. The more we know about how comprehension works, the more adequate our assumptions about its role in the generation of testimonial knowledge. Let me illustrate it with just one (however, quite prominent) example. According to a *local reductionist* view of testimony (Fricker 1994, 1995), for a receiver’s testimonial beliefs to be justified, the receiver has to establish whether the source of the information is trustworthy regarding this particular information on this particular occasion. Crucially, Fricker admits that, for her view of testimony to be accurate, it has to be the case that “it is not intrinsic to the state of understanding an utterance that it compels the hearer towards belief in what she grasps as being asserted” (1994, 157). Kourken Michaelian (2010) suggests, that Gilbert’s studies put some pressure on this assumption, and thus on the whole of Fricker’s account. As I argued extensively in Section 3.2, although Gilbert’s view is a little too strong, we actually *are*

compelled to believe what we understand if only the information passes the gatekeeper of validation. Moreover, the piece of information which passes validation becomes a part of the informational background and increases the probability that further inputs coherent with it will pass validation as well. Therefore, the model offered in this chapter turns out to be highly relevant to the epistemology of testimony. I come back to this topic in Chapter 4.

Chapter 3

The justification of comprehension-based beliefs

Abstract. What justifies our beliefs about what other people say (henceforth, *comprehension-based beliefs*)? According to *epistemic inferentialism*, the justification of comprehension-based beliefs depends on the justification of other beliefs, e.g., beliefs about what words the speaker uttered or even what sounds they produced. According to *epistemic non-inferentialism*, the justification of comprehension-based beliefs does not depend on the justification of other beliefs. This chapter offers a new defense of epistemic non-inferentialism. First, I discuss three counterexamples to epistemic non-inferentialism provided recently by Brendan Balcerak Jackson (2019) (“Against the perceptual model of utterance comprehension”, *Philosophical Studies* 176:387–405). I argue that only one of Balcerak Jackson’s counterexamples is effective, and that it is effective against only one version of epistemic non-inferentialism, viz. *language comprehension dogmatism*. Second, I propose an alternative version of epistemic non-inferentialism, viz. *comprehension-process reliabilism*, which is immune to these

counterexamples. I conclude that we should follow Balcerak Jackson in his rejection of language comprehension dogmatism but not all the way to the endorsement of epistemic inferentialism. Comprehension-process reliabilism is superior to both these alternatives.

1 Introduction

My friend Ruth and I have a chat over coffee. Thanks to my ability to understand what Ruth says, during our conversation I acquire two types of beliefs. First, I acquire beliefs that Ruth said so and so, e.g., that Ruth said that it's full moon tonight. Let us call them *comprehension-based beliefs*. Second, I acquire beliefs about the world based on the things that Ruth asserted, e.g., that it's full moon tonight. Let us call the second type of beliefs *testimony-based beliefs*. This chapter is devoted to comprehension-based beliefs; more specifically, to the issue of their justification.

A view about comprehension-based beliefs which has attracted considerable attention in the recent philosophical literature is the so-called *perceptual model* (cf. Fricker, 2003; Pettit, 2010; Brogaard, 2018, 2019). It is relatively easy to identify the intuition behind the perceptual model: comprehension is in important respects similar to perception, and thus comprehension-based beliefs are somehow similar to perceptual beliefs. Unfortunately, it is much more difficult to spell out the details of the perceptual model. Different authors have different opinions regarding what exactly accounts for the similarity between comprehension and perception. Additionally, it is not always clear whether the perceptual model is a view about cognition (i.e., the view that the process of formation of comprehension-based beliefs is similar to the process of formation of perceptual beliefs), epistemology (i.e., the view that the structure of justification of comprehension-based beliefs is similar to this of perceptual beliefs), or both.

For these reasons, in my current discussion, I will abandon the label *perceptual model* in favor of more precise categories. When it comes to the epistemology of comprehension-based beliefs, the perceptual model is often contrasted with an inferential model. Therefore, I will use the name *epistemic inferentialism* for the view that justification of the comprehension-based beliefs is *non-immediate*, i.e., that it depends on the justification of other beliefs, e.g., beliefs about what words the speaker uttered or even what sounds they produced. Respectively, I will call *epistemic non-inferentialism* the view according to which justification of comprehension-based beliefs is *immediate* and *does not* depend on the justification of other beliefs.¹

In his recent paper, Brendan Balcerak Jackson (2019) offers a critique of what he calls *perceptual model*. In fact, he criticizes a version of epistemic non-inferentialism: *language comprehension dogmatism* (or simply *dogmatism*). He puts forward three counterexamples designed to show that dogmatism fails to explain our intuitions about the justification of comprehension-based beliefs. The same task, according to Balcerak Jackson, is easily fulfilled by epistemic inferentialism. Crucially, Balcerak Jackson claims that his counterexamples constitute an abductive argument for epistemic inferentialism:

Taken collectively, the cases present a very strong abductive argument for concluding that the hearer's justification for her belief about what is said typically depends on her justification for believing that the speaker has uttered a particular sentence. If so then her justification is not immediate, contrary to what the perceptual model claims. (2019, 402)

This argument has an important flaw (which, I believe, results at least in part from the fact that Balcerak Jackson falls prey to the ambiguity of the term *perceptual model*). In a nutshell, Balcerak Jackson mistakenly takes his arguments against *a particular version of* epistemic non-inferentialism (i.e.,

¹Cf. Pryor (2005), Goldman (2008).

language comprehension dogmatism) to provide an abductive support for the view that justification of comprehension-based beliefs is inferential.

In this chapter, I argue, first, that Balcerak Jackson's critique is far less damaging to dogmatism than he thinks. Only one of his three counterexamples poses a serious threat to the theory. Second, I present another version of epistemic non-inferentialism, *comprehension-process reliabilism*, which resists all the counterexamples. Therefore, despite what Balcerak Jackson suggests, his critique does not provide motivation for inferentialism. At most, it helps us choose between alternative versions of epistemic non-inferentialism.

I proceed as follows. In Section 2, I introduce language comprehension dogmatism. In Section 3, I present Balcerak Jackson's counterexamples and discuss to what extent are they effective against this view. In Section 4, I outline comprehension-process reliabilism: a version of epistemic non-inferentialism immune to Balcerak Jackson's counterexamples.

2 Language comprehension dogmatism

Balcerak Jackson presents Fricker's (2003) view as a prototypical non-inferentialist position. According to Fricker, language comprehension gives rise to conscious representational states of *quasi-perception*, which confer prima facie justification on beliefs about what other people say.² Balcerak Jackson quotes Fricker:

When a hearer quasi-perceives that someone is saying that p to her, this is how things seem to her objectively to be, no less... Taking her aural experience at face value, that is what she would judge to be the case... Thus a quasi-perception of the content

²Fricker's quasi-perceptions are similar to phenomenal seemings produced by visual perception. I will use *quasi-perceptions* and *seemings* interchangeably to refer to all kinds of conscious experiential states taken to be immediately justificatory for comprehension-based beliefs.

and force of a heard utterance is, by its intrinsic nature, a prima facie though defeasible ground for belief. (Fricker, 2003, 341)³

A more recent version of this view is developed by Berit Brogaard (2018). In her own words:

...our experiences of what is said immediately justify our beliefs about what is said without any reliance on further belief or theorizing, at least in the absence of defeaters. So, in the absence of defeaters, our beliefs about the meanings of utterances are justified. (Brogaard, 2018, 2969)

I will use *language comprehension dogmatism* as an umbrella term for both these theories.⁴ What is central to language comprehension dogmatism is the claim that quasi-perceptions provide an immediate prima facie justification for comprehension-based beliefs. It is *immediate* in that *it is not* based on the justification of any other beliefs.⁵ It is *prima facie* in the standard Pollockian sense (Pollock 1986) of being good unless defeated. There are two kinds of defeaters: *undercutting* defeaters target subject's grounds for believing in a given proposition, e.g., the evidential connection between a seeming and the relevant belief; *rebutting* defeaters, on the other hand, are reasons to believe in the negation of the proposition or some other proposition incompatible with it. For example, I would have an undercutting defeater for a perceptual belief that I see a sheep if I knew that I have just taken a drug that is likely to cause visual hallucinations. I would have a rebutting defeater for the same belief, if I were told by the owner of the animal that it

³Notably, according to Fricker, quasi-perceptions represent both the content and the illocutionary force of the comprehended utterance. This creates further complications that are way too heavyweight to be addressed in this essay. I will point out just one of them in Section 3.1, because it contaminates an element of Balcerak Jackson's argumentation.

⁴Although in the next section I will discuss some of their idiosyncrasies.

⁵Including the belief that one's language comprehension faculty is reliable (cf. Pryor, 2005; Steup, 2018).

is actually a dog, which, when its fur gets long, looks very much like a sheep. According to dogmatism, unless a hearer⁶ is in possession of a defeater of one or the other type, their belief that it was said that p is justified if it seems to them that what was said was that p . A reader familiar with the debate about perceptual justification will immediately recognize similarities to Huemer's *phenomenal conservatism*:

Phenomenal conservatism: If it seems to S as if p , then, in the absence of defeaters, S thereby has at least some degree of justification for believing that p . (Huemer, 2007, 30)

Just like phenomenal conservatism, language comprehension dogmatism is a non-doxastic theory of justification, i.e., it rejects the claim that justification of comprehension-based beliefs has to depend on justification of other beliefs. It is, however, an *evidentialist* theory; seeming states provide evidential support for beliefs about what a speaker said. Thus, in both phenomenal conservatism and language comprehension dogmatism beliefs are, at least *prima facie*, evidentially justified by non-doxastic states.

Before moving forward, let us briefly examine whether seemings of what was said are in fact non-doxastic states. A fairly standard way of deciding whether a state is non-doxastic is by testing it against the background of a known illusion. Even if we know that the lines in the Müller-Lyer optical illusion are equally long (e.g., we examined them using a precise ruler) they still seem to differ in length. The seeming does not go away when confronted with a strongly justified belief that the lines are equally long. Assuming that it is rational to modify a false belief in the light of such glaring counter-evidence, we conclude that the representational state presenting the lines as unequal is not a belief. Rather, it is a non-doxastic seeming.

⁶A broader story about language comprehension will include also comprehension of written text and, plausibly, other types of linguistic inputs (e.g., sign language). For simplicity, in this essay, I will follow Balcerak Jackson in focusing on comprehension of spoken utterances and thus use *hearer* instead of the more medium-neutral *receiver*.

Can we think about a parallel situation in language comprehension? Brogaard (2018, 2975-2976) gives a great example. On the internet one can find tons of videos of cats making noises that sound as if they were saying “I love you!” None of us, except maybe a few over-the-top cat enthusiasts, believe that cats can say “I love you!”⁷ Yet the seeming that they do so remains immune to our belief. Therefore, seemings of what was said pass the test for being non-doxastic. Given that epistemic inferentialism claims that *all* comprehension-based beliefs are justified by inferential relations between beliefs, dogmatism is an obvious example of epistemic non-inferentialism.

3 Balcerak Jackson’s abductive argument for epistemic inferentialism

As I mentioned in the introduction, Balcerak Jackson offers three counterexamples to dogmatism, which he takes to constitute an abductive argument for epistemic inferentialism. Let us first see whether Balcerak Jackson’s counterexamples undermine dogmatism. Each of them is designed to reveal a different flaw in this theory. The first one is supposed to show that dogmatism can only *assume*, but cannot really *explain* why and how the prima facie justification provided by quasi-perceptions can be defeated. The second is supposed to show that, in certain situations, our background beliefs might actually justify a belief that someone *must have uttered some other sentence* than the one whose meaning we have quasi-perceived. The point of the third counterexample is to show that in the clairvoyant-style cases language comprehension dogmatism wrongly predicts that a hearer is justified in believing that what was said was that *p*, as long as it seems to them that what was said was that *p*.

⁷Even though they can be *trained* to produce these sounds it still would not count as *saying*.

3.1 Rabbits, habits, and minimal pairs

It will be useful for the discussion to quote Balcerak Jackson’s counterexamples at length.

Rabbit Habit

Ordinary competent English hearer Hans hears speaker Sven produce what sounds to him like an utterance of the following sentence:

(2) I have too many bad rabbits.

As it happens, Hans is being treated for a certain illness, and his doctor has warned him that the medication he is taking frequently has the curious side effect that patients regularly confuse [h] and [r] sounds in others’ speech. This means, for example, that for someone taking the medication there is a high likelihood that an utterance of ‘habit’ will sound like an utterance of ‘rabbit’ and vice versa. (Balcerak Jackson, 2019, 391)

Intuitively Hans is not justified in believing that Sven has said that he has too many bad rabbits. But why? An answer available to an epistemic inferentialist is that: (i) Hans has a justified belief that he is likely to misperceive [h] and [r], which in turn (ii) justifies his belief that Sven did not *utter the word* “rabbit” at the end of his utterance of (2), which finally (iii) defeats the justification for Hans’s belief that Sven *said* that he has too many bad rabbits. Obviously, a supporter of epistemic non-inferentialism cannot appeal to such inferential links between different doxastic states.

It is stipulated in Rabbit Habit that Hans has a quasi-perception of what was said in (2). According to dogmatism, this quasi-perception is *prima facie* justificatory for his belief about what Sven said. Since it is justified only *prima facie*, it might be defeated. However, Balcerak Jackson argues: “What the perceptual model [dogmatism] needs to explain is how the information

that Hans has from his doctor manages to defeat the prima facie justification provided by his quasi-perception.” (2019, 392). According to him, dogmatism does not offer such an *explanation*.

It would seem that there is a simple and natural answer available to a dogmatist. His doctor’s warning is a reason for Hans to doubt the reliability of his language comprehension faculty, and thus it undercuts the evidential support between the seeming that Sven said so and so and Hans’s comprehension-based belief—it serves as an undercutting defeater. However, Balcerak Jackson is not satisfied with such an answer. Let us see why.

Balcerak Jackson invites us to compare Rabbit Habit with a case in which the doctor says that the medication Hans is taking is likely to affect his perception of voice pitch or accent. If the warning in Rabbit Habit undermined Hans’s trust in the reliability of his language comprehension faculty, this one should as well. Yet, according to Balcerak Jackson, if the doctor’s warning concerned voice pitch or accent, we would still be inclined to take Hans to be justified in his beliefs about what is said: “*Neither of these warnings would have had any effect at all* on Hans’s justification for his belief about what was said.” (Balcerak Jackson, 2019, 392, *emphasis mine*).

This, I take it, is the core of Balcerak Jackson’s challenge. According to him, it is not enough to explain why a warning about possible misperception of the speaker’s utterance defeats the prima facie justification of comprehension-based beliefs. A dogmatist has to explain why some such warnings (about the confusion of phonemes) do, while others (about misperception of voice pitch or accent) do not.

I think that Balcerak Jackson is simply mistaken in assuming that a warning about misperception of voice pitch or accent would not defeat the justification of comprehension-based beliefs. In fact, it would have the exact same effect on Hans’s justification as the warning about possible confusion of phonemes; it would be an undercutting defeater. To see this let us first analyze what is so special about Rabbit Habit.

What is special about Rabbit Habit is that it is built around a minimal pair. English words “rabbit” and “habit”, just like “pin” and “bin”, “pen” and “pan,” constitute a minimal pair, i.e., a pair of words which differ in only one phonological element. In consequence, a misperception of phoneme [h] as [r] results not only in that it does not seem to Hans that Sven said “I have too many bad habits” but that *it seems to him* that Sven said “I have too many bad rabbits.” Moreover, “rabbit” and “habit” are of the same grammatical category. “I have too many bad rabbits,” even if potentially unexpected in the context of the conversation, is a grammatically correct English sentence.

Imagine now that what Sven says is “I am happy” or “It starts to rot.” I guess that Balcerak Jackson would not be inclined to argue that it is a challenge for a dogmatist to explain why Hans is not justified in believing that Sven said “I am reppy” or “It starts to hot,” respectively. Therefore, it is not the misperception of phonemes *per se* that makes Rabbit Habit tricky, but the fact that “rabbit” and “habit” constitute a minimal pair and that (2) is a grammatically correct English sentence.⁸

Balcerak Jackson argues that the warning about misperception of voice pitch or accent would not have *any effect* on the justification of comprehension-based beliefs. But we can build similar “minimal pair” examples around misperception of voice pitch and accent. First, intonation (i.e., the variation of pitch) often works in spoken English as a force marker. The rising intonation marks asking or interrogating: “You just walked in without ↗ knocking?”, while the falling intonation marks assertion: “You just walked in without ↘ knocking.” Since Balcerak Jackson follows Fricker (2003) in characterizing quasi-perceptions and comprehension-based beliefs as representing both the content and the force of heard utterances, we can think about a scenario in which Hans, informed by his doctor that he might be misperceiving voice

⁸An additional weakness of Balcerak Jackson’s case is that there is another occurrence of a “suspect” phoneme ([h]) in “I have too many bad habits.” How can Hans be sure that Sven did not want to say “I rav too many bad rabbits” thinking (mistakenly!) that the word “rav” (rhymes with “have”) means to kill in rage?

pitch, has a seeming as of Sven asserting “You just walked in without knocking.” In this case, contrary to what Balcerak Jackson assumes, Hans would not be justified in believing that Sven asserted that Hans has just walked in without knocking.

Similarly, with the misperception of accent. For example, a hearer who misperceives “special” as “spatial” due to a drug-induced failure “to compensate for the tensing of lax vowels characteristic of speech in southeast Ohio” (Bond, 2005, 298) would not be justified in believing that the speaker said “I have a spatial task for you.”

To sum up. We can build “minimal pair” cases similar to Rabbit Habit involving the misperception of voice pitch and accent. Claiming that the misperception of phonemes is somehow special, Balcerak Jackson guided us to a long detour. But we are back to where we started. All that is important from the perspective of the dogmatist theory of defeat is that the hearer is warned about the unreliability of their language comprehension faculty, and thus they have an undercutting defeater for their comprehension-based beliefs. If one has such a defeater, it does not matter that they have a seeming as if the speaker said something meaningful, e.g., “I have too many bad rabbits.” The defeater undercuts the evidential support between the seeming and the comprehension-based belief. Thus, dogmatism *can* explain why Hans is not justified in believing that Sven said “I have too many bad rabbits.”

3.2 He must have said “coat”

Let us now take a look at the second counterexample offered by Balcerak Jackson.

New Goat

Helena and Sven are in conversation, when it seems to Helena that Sven utters the following:

(4) I just bought a new goat.

Helena is initially puzzled by Sven's utterance. It seems to her that Sven has said that he just bought a new goat. But this would be a bizarre thing for him to say. (He hates goats, and he lives in a small city apartment that doesn't allow pets.) But Helena has every indication that Sven is speaking sincerely and honestly, and not ironically or metaphorically. All in all, Helena has very strong reason to believe that Sven did not, in fact, say that he had just bought a new goat, and so she concludes that she must have misheard him; perhaps she mistook an utterance of 'coat' for an utterance of 'goat'. (Balcerak Jackson, 2019, 400)

Crucially, Balcerak Jackson stipulates, that in this case, Helena has a strong defeater for the belief that Sven said that he bought a new goat: "if it is not already plausible from the above description, the case can be modified as necessary. (Perhaps the conversation takes place in front of dozens of attentive and reliable witnesses, all of whom sincerely assure Helena that Sven did not say that he had bought a new goat.)" (Balcerak Jackson, 2019, 400). Therefore, the task is *not* to explain why her belief that Sven said that he just bought a new goat is defeated. The actual task is to explain why "in this case there is another belief for which Helena *does* have justification, namely the belief that Sven must have uttered some other sentence than (4)." (Balcerak Jackson, 2019, 400; *emphasis in the original*).

Contrary to Balcerak Jackson, I do not take New Goat to be problematic for dogmatism (or any other version of epistemic non-inferentialism, for that matter). Obviously, epistemic non-inferentialism *does not* claim that *every* justified belief about what other people say is justified non-inferentially.⁹ The claim is that comprehension-based beliefs of a competent language speaker,

⁹For example, imagine that my friend is on the phone with his girlfriend and I hear him say "How can you say I'm needy?" I form a justified belief that my friend's girlfriend said that he is needy. Justification of this belief is obviously inferential.

formed through the exercise of their language comprehension faculty, are *typically* justified immediately. In New Goat, Helena’s belief that Sven said that he just bought a new goat was *prima facie* immediately justified by her seeming. However, as stipulated, the justification was rebutted. Believing that Sven *did not* say that he just bought a new goat (but curious about what he said), Helena further *inferred* that he must have said something that sounds similar but makes more sense in the present context. She concluded that he must have said “I just bought a new coat.” A supporter of epistemic non-inferentialism can readily agree that justification for the belief about *what Sven must have said is* non-immediate, i.e., inferential.

What is interesting about this case is the “initial puzzlement” that Helena experiences upon hearing Sven’s utterance. Of course, this is just an element of Balcerak Jackson’s description of the New Goat scenario; nevertheless, I think that this description is quite realistic and highlights a noteworthy phenomenon. To explain it, we should appeal to the fact that language comprehension is a predictive process (see, e.g., Kuperberg and Jaeger, 2016).¹⁰ Given the context of the conversation and other things she knew about Sven, Helena’s confrontation with the stimulus that her comprehension system identified as “goat” at the end of Sven’s utterance, triggered strong surprisal, which she might have experienced as puzzlement. Plausibly, this puzzlement prompted attention relocation in effect of which Helena started consciously analyzing what Sven has actually said. Eventually, she inferred that he must have uttered some other word than “goat”, probably: “coat”. Notice that in this last inference Helena must have appealed to her expertise with English minimal pairs (see section 3.1). Otherwise, why wasn’t her first guess that Sven must have said “fridge”?

3.3 Seemings out of the blue

Let us now move to the last counterexample offered by Balcerak Jackson.

¹⁰See also my discussion in Section 3.2 of Chapter 1 of this thesis.

Interpretive Clairvoyance

Hans is meeting Sophia for the first time at an informal gathering of United Nations workers. He asks Sophia about her area in an attempt to make small talk, but when she answers, it sounds to Hans exactly as if Sophia is speaking some foreign language that is entirely unknown to him; the sounds she is making strike Hans as nothing more than an unbroken stream of unfamiliar vocal noises. And yet at the same time, Hans has a quasi-perception as of Sophia saying that she is an economist. (Balcerak Jackson, 2019, 396)

As far as the intuitions go, Hans is not justified in believing that Sophia said that she is an economist. His quasi-perception of the content of Sophia's utterance came out of the blue. And yet, according to dogmatism, the sole fact that Hans has this quasi-perception provides justification for his belief "and none of the information at his disposal seems to do anything to call that into question." (Balcerak Jackson, 2019, 397).

I agree that Interpretative Clairvoyance is an effective counterargument to a version of dogmatism according to which *prima facie* justification depends *solely* on the possession of the relevant seemings or quasi-perceptions. But does it work against more restrictive versions of dogmatism? In the remainder of this section, I will take a look at a few such proposals.

First, let us examine Brogaard's *sensible dogmatism* (2013, 2018). According to Brogaard, not all seemings that *p* confer *prima facie* justification on beliefs that *p*. Only seemings that are "grounded in the content of [subject's] perceptual, introspective, or memory-related experience" (Brogaard 2013, 278) do so. In what sense are seemings *grounded* in the content of experience?

Content Grounding: A seeming of the form [It seems to A as if q] is grounded in a content *p* of a particular perceptual, introspective, or memory-related experience *e* had by A iff [Reliably(if *p* is

a content of e , then it seems to A as if q) and Reliably(if it seems to A as if q , then q)]. (Brogaard, 2013, 277)

The intended reading of *reliably* here is “in the majority of close hypothetical situations.” Thus, a seeming as if q is grounded in a content p of an experience e if and only if in the majority of close hypothetical situations experiences with the content p trigger seemings as if q , and in the majority of close hypothetical situations the world is as it seems.

I think that sensible dogmatism does not suffice to avoid Interpretive Clairvoyance. It is already stipulated that what Sophia says sounds to Hans like some foreign language. But according to sensible dogmatism, whenever Hans has *some* perceptual experience e that fulfills the condition for Content Grounding, his seemings as if Sophia said so and so should confer justification on his comprehension-based beliefs.¹¹ Imagine, for example, that every time Hans hears Sophia speak he has a sensation of a piano melody (different for different utterances) accompanied with a quasi-perception of the utterance meaning. In the majority of close hypothetical situations, experiences of a particular piano melody trigger the same seemings. If Hans has a sensation of melody a , he has a seeming as of Sophia saying that she is an economist, if he has a sensation of melody b , he has a seeming as of Sophia saying that it is her first time at the gathering of United Nations workers, etc. Moreover, in the majority of close hypothetical situations, if it seems to Hans that Sophia said that p , Sophia, in fact, said that p . Still, we are not inclined to say that Hans’ beliefs about what Sophia says are justified, even if there is a *systematic* relation between the piano melodies he experiences and his accompanying seemings.

Fricker’s (2003) theory is also a restrictive version of language comprehension dogmatism. According to Fricker’s definition of occurrent utterance understanding:

¹¹For a similar argument against sensible dogmatism see (Lyons 2015a).

To understand a particular heard utterance of [a sentence] S, U(S), occurrently is to enjoy a correct quasi-perceptual representation of the content and force of U(S), and to do so in part as the result of one's dispositional understanding of S—that is, through the exercise of *a stably possessed, internally constituted reliable capacity* to enjoy such representations with respect to utterances of S. (Fricker, 2003, 346, *emphasis mine*)

However, this is still not enough to deal with Interpretative Clairvoyance. Arguably, not *any* stably possessed, internally constituted reliable capacity to enjoy such representations would do. We would not say that my comprehension-based beliefs are prima facie justified by quasi-perceptions of utterance meaning, if these quasi-perceptions were triggered by sensations of piano music, smells, or colorful afterimages.

Here is one more potential line of defense for language comprehension dogmatism. A dogmatist could say that perceptual beliefs are prima facie justified by seemings that are not only grounded in experience (as it is characterized by Brogaard) but grounded in an *appropriate experience*. In the case of language comprehension, the appropriate experience might be, for example, the experience of words in a given language. Not just any experience of words in a given language would do, however. A dogmatist would have to stipulate that seemings of utterance meaning are grounded in the experience of *the appropriate* words. But how can we establish which words are appropriate to trigger a given seeming? Obviously, we cannot simply say that these are the words which have *the same meaning* which the hearer quasi-perceives, because the utterance meaning and word meaning are two very different things.¹² The case remains open.

To sum up. I argued that dogmatism has all the tools necessary to answer Rabbit Habit and New Goat. Nevertheless, I agree with Balcerak Jackson that neither of the existing versions of dogmatism successfully deals with

¹²Cf. Drożdżowicz (2019).

Interpretative Clairvoyance. I am not claiming that it is impossible to develop a version of dogmatism immune to this counterexample, but I will leave this task to devoted supporters of dogmatism. In the next section, I will argue that even if there is no hope for language comprehension dogmatism, it is too early to declare the victory of epistemic inferentialism. Balcerak Jackson’s abductive argument does not go through, because there is another non-inferential theory of justification of comprehension-based beliefs which deals with all counterexamples discussed in this section.

4 Comprehension-process reliabilism

Even though Balcerak Jackson’s discussion focuses on language comprehension dogmatism, he is ultimately interested in establishing that “the hearer’s justification for her belief about what is said is not immediate” (Balcerak Jackson, 2019, 402). Discarding dogmatism would be sufficient to achieve this goal if dogmatism was the only available theory of immediate justification of comprehension-based beliefs. But it is not. In this section, I outline another such theory, i.e., *comprehension-process reliabilism*.¹³

According to *process reliabilism*, a belief that p of a subject S is prima facie justified if it is generated by a reliable belief-forming process (cf. Goldman, 1979, Goldman and Beddor, 2016). Comprehension-process reliabilism is simply an application of process reliabilism to a particular type of belief-

¹³Discussing Interpretative Clairvoyance, Balcerak Jackson considers the possibility of rescuing the perceptual view by mixing it with reliabilism but does not find it promising. He claims that it is “the essence of the perceptual model that taking one’s quasi-perception at face value *is* a process that can confer justification” (Balcerak Jackson, 2019, 398), and therefore a supporter of perceptual model cannot say that Hans’s comprehension process in Interpretative Clairvoyance does not confer justification on his comprehension-based beliefs. Here again, Balcerak Jackson falls prey to the ambiguity of the term *perceptual model*: does it refer to dogmatism or epistemic non-inferentialism? Taking quasi-perceptions at face value is “the essence” of dogmatism, i.e., the evidentialist version of epistemic non-inferentialism. Below I will provide a non-evidentialist process reliabilism for comprehension-based beliefs. If we take *perceptual* to mean epistemically non-inferential, my account counts as a version of perceptual model.

forming process, namely, language comprehension:

Comprehension-process reliabilism: (i) a belief that p is prima facie justified if it is the result of a reliable belief-forming process; (ii) language comprehension is a non-inferential, reliable belief-forming process.¹⁴

This is a general recipe. To turn it into an actual account, one has to say more about the process of language comprehension and why it is reliable.

Dean Pettit (2010) comes close to formulating a version of comprehension-process reliabilism. First, he criticizes epistemic inferentialism as well as dogmatism, which suggests that he wants to defend a different kind of epistemic non-inferentialism.¹⁵ Second, according to Pettit, linguistic competence is *warrant conferring* but not *warrant apt*.¹⁶ It is not warrant apt because it does not stand in need of any further evidence. It is warrant conferring in that beliefs based on the exercise of linguistic competence are justified.¹⁷

Pettit provides reasons why linguistic competence is not warrant apt. He claims that it is deployed sub-personally and (at least to some extent) in a modular way. The question of warrant aptness simply does not arise with respect to such faculties. Just as our visual system, language comprehension does not require further evidence to confer justification to the beliefs it

¹⁴Given my goal here, I stipulate in the definition that language comprehension is *non-inferential* and, therefore, commit myself to epistemic non-inferentialism. One could drop this condition and formulate an inferentialist version of comprehension-process reliabilism. (Thanks to Derek Ball for helpful discussion here.)

¹⁵Balcerak Jackson mentions Pettit among representatives of the perceptual model (2019, 389, footnote 4), which is yet another indication that he *does not* restrict the perceptual model to language comprehension dogmatism.

¹⁶For the purposes of the current discussion, I will use warrant interchangeably with justification.

¹⁷It is worth highlighting that Pettit (2010) focuses on *linguistic competence* (competence with syntax and semantics of a given language) and not the process of *language comprehension* (exercise of this competence). However, linguistic competence is warrant conferring only insofar as language comprehension produces (at least prima facie) justified beliefs.

produces. In the absence of counter-evidence suggesting that it is *not* working properly, we are justified in taking its outputs at face value. But, as he focuses mostly on the critique of epistemic inferentialism, Pettit does not develop his alternative account in any detail; in particular, he does not really explain *why* the exercise of linguistic competence is warrant conferring, and thus, why we should take comprehension-based beliefs to be prima facie justified.

I will now outline an account which *does explain* why the process of language comprehension is warrant conferring. It draws on Peter Graham's (2010) teleological process reliabilism for testimony-based beliefs.¹⁸ Just as a reminder: while comprehension-based beliefs are beliefs of the form *that S said that p*, e.g., that Ruth said that it's full moon tonight; testimony-based beliefs are beliefs about the world, e.g., that it's full moon tonight, obtained through testimony.

Since Graham's target are testimony-based beliefs he does not focus on the process of language comprehension *simpliciter* but on what he calls *comprehension-with-filtering*. *Filtering* may be characterized as a suite of cognitive counter-deception mechanisms that prevent receivers of testimony from being easily manipulated by testifiers.¹⁹ Graham enumerates a list of empirical research supporting the claim that, starting at a very young age

¹⁸There are, obviously, other paths one could follow to develop their favorite version of comprehension-process reliabilism. One of them is suggested by Jack Lyons (2009). According to Lyons, beliefs are prima facie justified if they are produced by processes rooted in so called *primal systems*. Primal systems are, in turn, characterized as systems: (i) whose inner workings are not accessible to introspection; (ii) which result from an innate state of an organism developed by learning. They are, thus, quite similar to Fodor's modules (Fodor, 1983; for a detailed comparison with Fodor's view, see Lyons, 2015b). The most widely discussed primal system is perception, but Lyons suggests: "Some beliefs about the speech of others are clearly perceptual. My auditory belief that so-and-so just said 'The cat is on the mat' is a straightforwardly perceptual belief. My belief that so-and-so just said that the cat is on the mat will be a perceptual belief only if the language comprehension system counts as a perceptual system; it very well might on the present understanding of perceptual systems, and I have no problem with this result." (2009, 135).

¹⁹This notion of *filtering* is very close to Sperber et al.'s (2010) *epistemic vigilance*; see my discussion in Chapter 4.

(some experiments involve 16-month-olds), we develop special attentiveness towards speakers who make obvious mistakes, do not sound confident, have been unreliable in the past, etc. Notably, according to Graham, filtering does not condemn us to epistemic inferentialism.

Filtering need not involve explicit awareness—belief or judgement—that counter-considerations are absent. Filtering need not involve reasons and reasoning. . . filtering involves sensitivity to counter-considerations: where there are counter-considerations of a certain sort, acceptance would be less likely. (Graham, 2010, 152)

As I have mentioned, Graham’s account is teleological. It is spelled out in terms of *etiological function*, *normal functioning*, and *normal conditions* (Millikan 1984). In a nutshell, an *etiological function* of any mechanism or feature is the effect of its ancestors, which explains why the mechanism or feature persisted. The etiological function of our hearts is to pump blood, because pumping blood was the effect that hearts were selected for, i.e., it is the effect of ancestors of our hearts which explains why we still have hearts. *Normal functioning* is working in the way that contributed to stabilizing the etiological function (e.g., pumping blood, in the case of the heart), and *normal conditions* are conditions sufficiently similar to those in which “that function has historically been performed” (Millikan, 1984, 34).

Here is how Graham applies the teleological apparatus to the problem of epistemic entitlement.²⁰ According to him, “entitlement attaches to beliefs in virtue of the normal functioning of the belief-forming process when the

²⁰Graham appeals to Burge’s (1993, 1996) distinction into two types of *warrant*: *justification*, which involves reasons and evidence a subject can often cite; and *entitlement*, which does not involve them. Given that, according to Graham, the positive epistemic standing of testimony-based beliefs does not depend on reasons or evidence, he says that testimony-based beliefs enjoy prima facie entitlement. In the present discussion, I will use *justification* and *entitlement* interchangeably. However, the reader should bear in mind that justification in the sense in which I am using it *does not* require reasons or evidence.

process has forming true beliefs reliably as an etiological function” (Graham, 2010, 156). Perception-based beliefs enjoy prima facie entitlement because the perceptual process has forming true beliefs reliably as a function. Similarly, Graham argues, comprehension-with-filtering has inducing true testimony-based beliefs reliably as its function. The reason why we are equipped with the cognitive mechanism of comprehension-with-filtering is that an ancestor of this mechanism was beneficial to our ancestors. More specifically, it allowed them to reliably acquire true beliefs about the world based on testimonies of their interlocutors.²¹

Since my goal in the present chapter is to provide an account of justification of comprehension-based beliefs, i.e., beliefs of the form *that the speaker said that p*, I do not need to engage here in the debate about the nature and epistemic status of filtering.²² Therefore, I would like to suggest a restriction of Graham’s theory to comprehension *simpliciter*.

Teleological comprehension-process reliabilism (TCR): (i) a belief is prima facie justified if it is based on a normally functioning belief-forming process that has forming true beliefs reliably as a function; (ii) language comprehension is a non-inferential, belief-forming process that has forming true comprehension-based beliefs reliably as a function.

I find it extremely plausible that forming true comprehension-based beliefs is the etiological function of the process of language comprehension. Hearts persisted because of their beneficial effect of pumping blood. Language comprehension persisted because of its beneficial effect of generating

²¹Graham makes a further, Millikan-style argument that it is an etiological function of *assertion* to induce true beliefs in hearers. Otherwise, assertion would not persist. “So unless hearers get something out of accepting reports, they will not accept them. And if they will not accept them, speakers will not benefit from making them. Then they will not get made.” (Graham, 2010, 160). In result, the function of assertion and the function of comprehension-with-filtering taken together support the prima facie justification of testimony-based beliefs.

²²I discuss it extensively in the next chapter.

true beliefs of the form *that S said that p*. That obtaining such true beliefs is beneficial to members of our species is rather uncontroversial; it contributes to successful communication and allows us to learn from and cooperate with others.

Finally, the restriction of Graham's teleological account to comprehension *simpliciter*, and thus to the generation of comprehension-based beliefs, seems to be necessary anyways. Forming true testimony-based beliefs is the function of comprehension-with-filtering of *assertoric utterances*. But what about other speech acts? Understanding that Ruth asked whether it's full moon tonight does not (at least not directly) generate the belief that it's full moon tonight; understanding that Ruth ordered to stop the car does not (at least not directly) generate the belief that I will stop the car, etc. A complete theory of language comprehension, applicable to all speech acts, seems to require the two-level structure with (i) comprehension-based beliefs on one level, and (ii) testimony-based beliefs for assertions, and their counterparts for other illocutionary forces, on the other.²³

Let us now see whether TCR can deal with Balcerak Jackson's counterexamples. Discussing Rabbit Habit in Section 3.1, I argued that it does not pose an actual problem to dogmatism. According to dogmatism, seemings provide *prima facie* justification to comprehension-based beliefs. As *prima facie*, this justification can be defeated, e.g., if the hearer has a reason to doubt that their language comprehension faculty is reliable. Thus, the warning about the possible misperception of phonemes, which Hans receives from his doctor, plays the role of a defeater undercutting the evidential support between his seeming and the respective belief. An equivalent answer is available to TCR. TCR is also a theory of *prima facie* justification and thus it leaves room for defeat. If, as it is in Rabbit Habit, a hearer *knows* that his language comprehension faculty is not functioning normally, the *prima facie*

²³The two-level structure is a theoretical assumption. I am not claiming that it has psychological reality, e.g., that when we comprehend an assertion, we have to go through both these levels.

justification of beliefs formed by this faculty is defeated, and the hearer is not justified in believing that the speaker said so and so.²⁴

What about New Goat? Discussing this counterexample in section 3.2, I have argued that it does not pose a problem for epistemic non-inferentialism. It is stipulated by Balcerak Jackson that Helena's justification for the belief that Sven said "I just bought a new goat" is defeated. It might be the case that Helena will end up forming another belief, i.e., the belief that Sven must have said "I just bought a new coat," but a supporter of any version of epistemic non-inferentialism (be it dogmatism or TCR) can readily agree that Hans's justification for this belief is inferential. Clearly, not all beliefs about what other people say (or must have said) are justified immediately. The belief that Sven must have said that he bought a new coat is not.

Finally, what about Interpretative Clairvoyance? Just as a reminder, the challenge posed by Interpretive Clairvoyance was to explain why Hans is not justified in believing that Sophia said that she is an economist, despite the fact that it seems to him that she said this exact thing. What is crucial for TCR is that it does not take justification to be grounded in experiential states (or in any other state, for that matter). It is not only a *non-doxastic* but a fully *anti-evidentialist* position. It might very well be the case that typically the language comprehension process is *associated with* certain phenomenal states, but the justification of comprehension-based beliefs is not *grounded* in these states. Comprehension-based beliefs are justified only as far as they originate from a normally functioning language comprehension process. Of course, Hans's belief in Interpretive Clairvoyance is not formed by a normally

²⁴There is a long-standing debate in epistemology, one in which I will not engage here, whether process reliabilism can offer a satisfactory account of epistemic defeat. The most popular strategy of explaining defeat in process reliabilism, Alvin Goldman's *Alternate Reliable Process (ARP) theory* (Goldman 1979; cf. Lyons, 2009, 2016), has been passionately discussed over the years (see Beddor, 2015, *forthcoming* for good overviews). For recent defenses of improved versions of ARP see Grundmann (2009) and Graham and Lyons (*forthcoming*). Here, I will simply assume that there is a theory of defeat available to process reliabilism.

functioning comprehension-process or any other process that has forming true-beliefs reliably as a function. Hans's ancestors *did not* have the faculty of interpretive clairvoyance, so this faculty does not have an etiological function at all. This is enough to explain why interpretive clairvoyance does not produce prima facie justified beliefs.

The above is not an all-out defense of comprehension-process reliabilism or TCR. Nevertheless, by outlining and making an initial case for this theory, I have shown that we are not doomed to epistemic inferentialism, even if Interpretive Clairvoyance defeats language comprehension dogmatism.

5 Concluding remarks

The goal of this chapter was to offer a new defense of epistemic non-inferentialism about comprehension-based beliefs, i.e., the view that justification of comprehension-based beliefs is immediate. I started by examining Balcerak Jackson's critique of a version of epistemic non-inferentialism, viz. language comprehension dogmatism. Only one out of three counterarguments provided by Balcerak Jackson turned out to be successful, but it was enough to reject dogmatism (at least until an improved version of this view is provided). Nevertheless, the rejection of dogmatism does not suffice to establish epistemic inferentialism. In the last section, I offered another version of epistemic non-inferentialism, viz. teleological comprehension-process reliabilism. I argued that this view successfully deals with all counterexamples provided by Balcerak Jackson.

Chapter 4

Effective filtering

Abstract. It is a popular view in social epistemology that we are equipped with a set of cognitive tools that help us to filter out unreliable testimony. But do we have reasons to assume that these tools are effective? In this chapter, I answer this question in two steps. In the first step, I argue that filtering is not real-time effective, i.e., the process of filtering, which takes place simultaneously with or right after the process of comprehension, does not allow a particular hearer on a particular occasion to prevent the formation of beliefs based on untrustworthy testimony. In the second step, I argue that filtering is long-term effective. The very fact that some hearers sometimes detect untrustworthy testimony puts pressure on speakers and increases the incentive for honesty. Finally, I argue that these considerations support a version of strong social anti-reductionism about testimonial entitlement.

1 Introduction

It is hard to imagine how limited one's knowledge about the world would be, had it not been for testimony. It is through testimony that I know that

Huascarán is in the Andes and that mitochondria contain DNA. It is also through testimony that I know where my parents spent their last holidays, and what dose of medicine works best for my dad's blood pressure. From trivia to vital information, our epistemic situation depends greatly on things we are being told.

A widespread opinion in the philosophical debate about testimony is that we are not gullible, i.e., we engage in some form of assessment of the trustworthiness of testimony we receive or, at least, we are sensitive to signs of its untrustworthiness.¹ This assumption plays a particularly important role in reductionist theories of testimonial warrant.² As stated in Elizabeth Fricker's classical paper:

...the hearer should be discriminating in her attitude to the speaker, in that she should be continually evaluating him for trustworthiness throughout their exchange, in the light of the evidence, or cues, available to her. This will be partly a matter of her being disposed to deploy background knowledge which is relevant, partly a matter of her monitoring the speaker for any tell-tale signs revealing likely untrustworthiness. (Fricker, 1994, 149-50)

This monitoring (or filtering)³ is a way of collecting reasons for speakers'

¹Cf. Shieber (2012, 2015).

²When I speak about *reductionism* in the epistemology of testimony, I refer to a set of views which tend to agree that the warrant of testimony-based beliefs depends on hearer's having some reductive (i.e., based not on testimony itself but, for example, perception, memory, or inductive inference) reasons for trusting their testifier. Proponents of reductionism include Hume (1975), Adler (1994), Audi (1997, 2004, 2006), Fricker (1987, 1994, 1995, 2006), Lipton (1998), Lyons (1997). *Anti-reductionists*, on the other hand, argue that testimony-based beliefs can be warranted without hearers' possessing such reductive reasons. Various forms of anti-reductionism have been defended by Reid (1970), Burge (1993, 1997), Coady (1973, 1992), Goldberg (2006, 2007, 2010, 2014), Goldman (1999), Graham (2006, 2010, 2015), Greco (2015), Simion (2020).

³Throughout the text I will use *monitoring* and *filtering* interchangeably. I will also use the noun *filter* to refer to a set of mechanisms, processes, or any cognitive tools that allow us to perform monitoring/filtering.

trustworthiness, and thus lies in the core of reductionism. However, even though reductionists tend to accuse anti-reductionists of painting hearers as helplessly gullible,⁴ there are ways of incorporating the monitoring requirement into anti-reductionist accounts.⁵ In his polemics with Fricker, Sanford Goldberg (2007) offers the following formulation of anti-reductionism:

A hearer H is epistemically justified in accepting (has the epistemic right to accept; is epistemically entitled to accept) another's testimony on occasion O so long as (i) there are no undefeated good (doxastic, factual, or normative) reasons not to accept the testimony, and (ii) on O H's acceptance was the outcome of a process that exhibited a 'counterfactual sensitivity' to the presence of defeaters (which, given (i), turns up no such defeaters on O). (2007, 168)

The monitoring requirement is captured in clause (ii) and takes the form of "counterfactual sensitivity." Goldberg (2007, 166) highlights that it is sufficient for hearers to *be on a lookout* for defeaters (as opposed to *looking for defeaters*) to avoid the charge of gullibility, and that this condition is compatible with anti-reductionism.⁶

Authors trying to carve up a middle ground between reductionism and anti-reductionism also suggest that, as recipients of testimony, we routinely monitor the trustworthiness of testifiers. To give just one example, Jennifer Lackey (2006) suggests that we do it by relying on "inductive evidence for believing. . . that reports made with sustained eye contact are typically sincere ones, or that reports made ably and confidently are typically competent ones." (p. 173).

This chapter takes up the question: "Is filtering effective?" My answer consists of two steps. In the first step (Section 2), I argue that filtering is

⁴Cf. Fricker (1994, 154; 1995, 404).

⁵Cf. Coady (1992, 47).

⁶Cf. Goldberg and Henderson (2006).

not real-time effective, i.e., filtering, which takes place simultaneously with or right after comprehension, does not allow a particular hearer on a particular occasion to prevent the formation of beliefs based on untrustworthy testimony. To establish this conclusion, I argue, first, that Michaelian (2010) and Shieber’s (2012, 2015) arguments against real-time effectiveness of filtering are immune to recent critique by Fricker (2016). Second, I outline and shortly defend an account of language comprehension, which suggests that filtering *cannot* be real-time effective.

Even if filtering is not real-time effective, it does not mean that it is not effective at all. As I argue in Section 3, filtering brings expected benefits only in the long run, i.e., that it is long-term effective. The very fact that some hearers sometimes detect untrustworthy testimony puts pressure on speakers and increases the incentive for honesty. Finally (Section 4), I argue that these considerations support a version of strong anti-reductionism about testimonial entitlement.

2 Filtering is not real-time effective

2.1 Michaelian, Shieber, and Fricker on the real-time effectiveness of filtering

Kourken Michaelian (2010) and Joseph Shieber (2012, 2015) are two authors who recently devoted a lot of attention to the problem of the real-time effectiveness of filtering.⁷ Both of them proceed from a compelling assumption: as long as we are interested in how filtering works and how impermeable it is, we cannot limit ourselves to theoretical considerations. Instead, we should consult relevant empirical research, mainly from the field of social psychology.

⁷Real-time effectiveness of filtering can be understood in terms of *reliability*: a filter is real-time effective if it takes place simultaneously with or right after the process of comprehension, and reliably prevents a particular hearer on a particular occasion from forming beliefs based on untrustworthy testimony.

Michaelian's (2010) paper is a polemic with Fricker's reductionist theory of testimonial justification and knowledge (Fricker 1987 and onward). For the purpose of his discussion, Michaelian breaks Fricker's theory down into two main components: (i) the claim that the reductive account to testimonial justification (reduction of testimonial justification to receiver's possession of non-testimonial reasons) is necessary, and (ii) the claim that it is available. The question of the effectiveness of filtering plays a crucial role in assessing the second claim. In Michaelian's reconstruction, the process of formation of testimonial beliefs postulated by Fricker:

...takes as input the subject's beliefs about what the speaker said, about the competence of the speaker, and about the honesty of the speaker and produces as output either a testimonial belief or no belief at all: if the subject believes that the speaker asserted that *P*, that the speaker was then competent with respect to *P*, and that the speaker was then honest with respect to *P*, the process outputs a belief that *P*; otherwise (e.g., if the subject believes that the speaker was incompetent with respect to *P*), the process outputs no belief. (Michaelian, 2010, 406).

Two features of this characterization of the belief-forming process will become central to my discussion in Section 2.2, but I would like to flag them here. Firstly, Fricker assumes that one of the inputs to the process (the one achieved by way of linguistic comprehension) is the receiver's belief about what a speaker said (Fricker, 1994, 150). Secondly, according to Fricker, the output of the process is either belief or no belief. As I hope to demonstrate in the next section, despite their popularity in the contemporary debate on testimony, both these assumptions are problematic.

For the time being, however, let us focus on what is crucial for Michaelian's critique of Fricker's theory. The reason why receivers are not gullible, according to Fricker, is that the process of formation of testimonial beliefs makes

use of their beliefs about honesty and competence of the testifiers. These beliefs, in turn, are acquired by means of—plausibly *unconscious* (Fricker, 1994, 150)—monitoring for signs of dishonesty and incompetence. What does this monitoring require? Fricker reassures us that it *does not* require “... an extensive piece of M15-type ‘vetting’ of any speaker before ... [the hearer] may accept anything ... [the speaker] says as true” (1994, 154). Moreover, Fricker optimistically assumes that “[e]xpert dissimulators among us being few, the insincerity of an utterance is very frequently betrayed in the speaker’s manner, and so is susceptible of detection” (Fricker, 1994, 150). But do we have good reasons for such optimism?

Michaelian argues that even though Fricker might be right that we frequently and casually monitor for competence and insincerity, there are no good reasons to assume that this monitoring is effective in the sense Fricker’s theory requires. Michaelian focuses on detection of a speaker’s dishonesty. He presents a line-up of empirical studies which suggest that our accuracy rates in deception detection are “only slightly better than fifty-fifty” (Levine et al., 1999, 126; cf. Bond and DePaulo, 2006; Levine, 2019b).⁸ There may be many reasons for such underwhelming results, but Michaelian suggests that it is neither that receivers do not monitor for cues to deception nor that there are no cues to deception to monitor for in the first place. Rather, the task of monitoring for cues to deception is very difficult: receivers plausibly do not monitor for all the cues to deception and the ones that they monitor for are subtle and vary significantly across agents and contexts (Vrij, 2000, 2004; Feeley and Young, 2000).

Shieber (2012, 2015) raises similar worries. He suggests that we can think about two types of monitoring strategies: the first, based on gathering positive evidence of the trustworthiness of a given testifier; and the second,

⁸The slightly better than chance accuracy (around 54%) might result from a *few transparent liars* effect (Levine, 2010). According to Levine’s metaphor, deception detection is similar to taking a test where you know answers to approx. 10% of the easiest questions, and guess all the rest.

based on sensitivity to indications of deception. Unfortunately, according to Shieber, social psychology does not give us reasons to assume that either of these strategies is reliable.⁹

When it comes to the first strategy, its success depends on a receiver's ability to detect cues to trustworthiness. However, according to Shieber, psychological literature suggests neither that there is a unified set of such cues nor that we are particularly good at spotting them. Instead, receivers tend to trick themselves into taking physically attractive (Chaiken, 1979) and generally likeable (Chaiken and Eagly, 1983) testifiers as trustworthy. Furthermore, sustained eye contact or "keeping straight face," commonly taken to be cues to trustworthiness, may be in fact negatively correlated with sincerity (Sitton and Griffin, 1981; Bond et al., 1985).

What about the second strategy? Shieber's overview of the psychological literature leads him to the same conclusion as the one formulated by Michaelian: it does not seem that our success rate in detecting deception is much better than chance (0.5). Moreover, Shieber quotes research according to which accuracy rates of professionals whose work may require sensitivity to deception, such as police officers, are not significantly better than those of laypeople (Kraut and Poe, 1980; DePaulo and Pfeifer, 1986; Köhnken, 1987). To sum up, we are bad at deception detection by default, and there is little room for improving this skill through training.^{10,11}

⁹Separately, Shieber discusses possible strategies for monitoring for the competence of a testifier, and reaches similarly pessimistic conclusions (cf. 2012, 9-12). For my current discussion, it is sufficient to focus on monitoring for a testifier's trustworthiness.

¹⁰The most optimistic conclusion drawn from existing empirical research on the improvement of deception detection accuracy (when it comes to real-time detection based on behavioral cues) is that "*certain* professions or *certain* subgroups within professions may develop particular sensitivity to *certain kinds* of lies" (O'Sullivan et al., 2009, 536, *emphasis mine*).

¹¹Some readers might be worried that Michaelian's and Shieber's arguments rely on an already outdated psychological literature. But this is not the case. The same conclusions about the real-time effectiveness of filtering are presented in recent meta-analyses and literature reviews (Hartwing and Bond, 2014; Hauch et al., 2016; Levine, 2019b; Solbu and Frank, 2019; Sternglanz et al., 2019). See also discussion below.

Fricker (2016) strongly resists the pessimism. What we have to take into account while assessing the effectiveness of filtering—Fricker suggests—is that that our communication takes place in different *testimonial environments* (T-environments). T-environments are individuated based on “what frequency and what manner of false testimony . . . [the receiver] might easily encounter” (Fricker, 2016, 96). For a filter to be effective is for it to be effective in a given T-environment; the same “narrowly” construed belief forming method might turn out to be reliable in one environment (e.g., one’s everyday T-environment full of family members, friends, and acquaintances), but not another (e.g., a T-environment full of habitual and expert liars).

Moreover, since it is well established that cues to deception differ across testifiers and contexts, if the empirical studies were supposed to help us assess the effectiveness of filtering, environments created in the experimental settings would have to be sufficiently similar to the everyday T-environments of participants. But this is not the case in the current empirical literature. Additionally, the effectiveness of the filter is not supposed to result solely from our ability to recognize perceptual cues to deception typically targeted by experimental studies. The other, and probably even more important component of the filtering toolkit, is the use of background knowledge about specific testifiers (including past records of their honesty; current motives, etc.) and general social context. Finally, even if there were empirical reasons to assume that we do not filter effectively, there are no reasons to assume that we cannot *learn* to filter effectively. According to Fricker, the jury is out on effectiveness of filtering until we demonstrate that: “Humans are *constitutionally incapable* of learning to respond discriminately to testimony.” (2016, 103, *emphasis mine*).

Fricker concludes that empirical data cited by Michaelian and Shieber is insufficient to vindicate the claim that filtering is not effective. This comes with no surprise, I think, given the requirements she imposes makes it virtually impossible to vindicate this claim using empirical methods. Fricker does

not explain in detail how finely we should individuate T-environments, however, given that T-environments are individuated by appeal to frequency and manner of false testimony one encounters, we can easily imagine that everyone would have their own everyday T-environment unlike the T-environment of any other person or even unlike their own T-environment at different times. If this sounds radical, let's take a look at the example Fricker provides.

ELLA: Ella, a teenager, has a circle of friends in which there is a strong norm of trust and honesty. They very rarely deceive each other—even to the point of preferring honesty to tactfulness. In the situation of Ella and her circle lying is infrequent, and anyone who lies is unpracticed, guilty and embarrassed. The liar shows tell-tale signs and Ella, no fool and perceptually well-attuned, is able to detect them. Then Ella's father gets a new job, and the family moves to a city in another part of the country. Ella goes to a new school with very different social conditions and mores. In her new school, lots of the pupils lie regularly in pursuit of their own selfish purposes and are practiced and proficient dissemblers. Ella retains the perceptual attunedness she previously possessed—her “narrow” T-method is the same. But in her new environment the old skill is not sufficient for the different and more taxing task of detecting when these cynical streetwise city kids, her new classmates, are lying. (Fricker, 2016, 97-8)

Now let's imagine that Ella takes part in an experimental study on deception detection. If it takes place before she moves to the new city, according to Fricker's standards, to actually measure Ella's detection deception skills, the setting of the study should replicate the environment of her old school. But what if the study took place a couple of months after she moved to the new school? With the environment of the new school becoming her new everyday T-environment, her filtering skills gradually attune to the reality of the life of

the “cynical streetwise city kids.” To measure her detection deception skills after spending a couple of months in the new environment, the study should replicate the T-environment of the new school. For now, we take into account still relatively broadly characterized T-environments of the old school and the new school but there is no principled reason why we should stop here. Every single student of each of these schools encounters a different set of testifiers (assuming that no one is their own testifier), which might affect the frequency and manner of the false testimony they are exposed to, thus, plausibly, every student has their own T-environment.

Reconstructing such specific individual differences is obviously impossible in an experimental setting. After all, it is the bread and butter of empirical research that an experimental design balances between keeping the environment as natural as possible while simultaneously controlling for variables that might affect the result. What comes close to fulfilling Fricker’s expectations are studies devoted to deception detection in intimate relationships: romantic relationships, friendships, or parent-child relationships (Mccornack and Parks, 1986; Evans et al., 2016; Levine and Knapp, 2018). Interestingly, many of these studies demonstrate that deception detection accuracy is actually lower in close relationships than between strangers (even though participants are more confident in their judgment). Others demonstrate that it is slightly higher in close relationships, but still no higher than 65% (Levine and Knapp, 2018).

Fricker openly discards the assumption that: “the data which show that recipients of testimony are bad at detecting lying concern studies in a very specific experimental setting; but the nature of the findings may nonetheless be such that it is likely that they will carry over to other situations in which testimony is given and received.” (Fricker, 2016, 102). Similar assumptions are fairly standard in social psychology. I do not want to imply that there are no problems with ecological validity of deception detection research, only that there is no reason to flat out dismiss them based on this observation.

What is lacking in Fricker’s critique, is an argument that real-time deception detection is uniquely environment-dependent and thus virtually impossible to investigate empirically.

I understand and sympathize with Fricker’s worry that, while appealing to empirical research in philosophical discussion, we might be tempted to cherry-pick studies which support our points. That is why it is important to, whenever it is available, take into account not only particular studies but also meta-analyses which allow us to recognize general tendencies and derive unified conclusions from bodies of research consisting of multiple studies. Meta-analyses available in deception-detection literature point consistently into the direction of only slightly better than chance accuracy of deception detection in general population and fairly limited possibility of its improvement by training (Bond and DePaulo, 2006; Hartwig and Bond, 2014; Hauch et al., 2016; Sternglanz et al., 2019).

In sum, I think that we should reject Fricker’s critique as based on unrealistic demands. Empirical research cited by Michaelian and Shieber (together with further research on deception detection published in the last ten years) make a very strong case against the real-time effectiveness of filtering. But I also think that we can do even more to demonstrate that filtering is not real-time effective. In the next section, I will take up Fricker’s challenge of demonstrating that: “Humans are constitutionally incapable of learning to respond discriminately to testimony.” (2016, 103).

2.2 Three models of language comprehension

At the beginning of the previous section, I presented a reconstruction of Fricker’s account of the testimonial belief-forming process and flagged out two of its features. First, besides the beliefs about the speaker’s honesty and competence, the belief-forming process takes as an input the receiver’s belief about what the speaker said (e.g., *that the speaker said that p*). Following the terminology introduced in Chapter 3, I will call such beliefs *comprehension-*

based beliefs. Second, based on the input it receives, the belief-forming process outputs either a testimonial belief that p or no belief at all.

Here is the model of language comprehension underlying Fricker's account: upon hearing or reading an utterance, receivers (i) form comprehension-based beliefs representing the speaker as asserting¹² certain content (p), and then (ii) either accept or reject p based on the assessment of the speaker's honesty and competence, which leads to the formation of a corresponding testimony based belief that p or no formation of belief. I will call this model *Optional Belief Formation (OBF)*. If one thinks, as most authors in the debate do, about filtering in terms of OBF, it is natural to assume that the filter is "located at the entrance" to our belief box and that its role is to keep contents of testimony from unreliable (dishonest or incompetent) sources from falling into the belief box. OBF is so popular across the current philosophical debate, that one could assume it is the only game in town. But it is not.¹³

In an alternative model (*Automatic Belief Formation (ABF)*), upon comprehending an utterance that p , a receiver automatically accepts the content of this utterance and forms a belief that p .¹⁴ Later on, they can reject the belief and remove *that* p from their belief box. However, rejection is an additional step, which requires extra time and cognitive resources. As I have already described in Chapter 2, viability of ABF was brought to the attention of philosophers by research conducted by psychologist Daniel Gilbert and his colleagues in the early nineties (Gilbert et al., 1990; Gilbert, 1991; Gilbert et al., 1993). In a series of experiments, Gilbert aimed to demonstrate that under cognitive and time pressure, people fail to reject comprehended contents that are explicitly identified as false. His studies suggest, *contra* OBF, that content-acceptance and belief-formation are not optional, but automatic and

¹²This is the case for assertoric speech. Plausibly, in cases of other speech acts, receivers represent speakers as asking, ordering, etc.

¹³It is worth mentioning that Michaelian (2010) is aware of this fact (p. 403, footnote 9). However, for the purpose of his discussion with Fricker, he assumes OBF to be true.

¹⁴Since (Gilbert et al., 1990), OBF and ABF are often called *Cartesian* and *Spinozan* model, respectively. Cf. Chapter 2 of this thesis.

mandatory. Gilbert’s research had a big impact on Ruth Millikan’s theory of language comprehension (Millikan, 2004, 2005) and Eric Mandelbaum’s (Mandelbaum, 2014; Mandelbaum and Quilty-Dunn, 2015) and Andy Egan’s (2008) theory of belief formation in general.¹⁵

If one thinks about filtering in terms of ABF, one has to assume that there is only filtering *ex post*. Nothing could prevent the initial formation of a belief that *p* based on comprehension of an assertion that *p*. This sounds (and indeed is) very radical. What about the comprehension of blatantly false and improbable assertions? Do I automatically form beliefs that Paris is in Germany or that the Earth is flat upon hearing these pieces of information being asserted?

As I argued extensively in Chapter 2, thirty years of empirical research on language comprehension since Gilbert’s seminal studies suggest that neither OBF nor ABF gives a fully accurate picture of language comprehension (Hasson et al., 2005; Street and Richardson, 2015; Kissine and Klein, 2013). Apparently, *contra* OBF, a lot of what we comprehend is in fact automatically accepted before the credibility of the speaker is taken into account (Spark and Rapp, 2011; Nadarevic and Erdfelder, 2013; Weil et al., 2019). However, *contra* ABF, we do not accept contents that are glaringly inconsistent with our active or easily accessible background knowledge. Besides the source-oriented filter commonly discussed in the philosophy of testimony, there is a content-oriented filter which the literature on language comprehension calls *validation*, and which prevents us from believing obviously false information (Singer, 2013, 2019; Schroeder et al., 2008; Richter et al., 2009; Isberner and Richter, 2014; Richter, 2015).¹⁶

Based on these considerations, in Chapter 2 I proposed a dual-stream model of language comprehension. The first stream is faster and entirely

¹⁵Mandelbaum (2014) offers an excellent discussion of Gilbert’s research and its philosophical consequences.

¹⁶The suggestion that filtering has to be both content-oriented and source-oriented appears also, e.g., in Sperber et al. (2010), Sperber (2013), Levine (2019a).

content-oriented. It updates contents of comprehended assertions into our belief box unless they are filtered out by validation. For example, if Liz comprehends Tom’s utterance that p , this stream processes only the content of the utterance (p), and forms a belief that p , unless p is obviously false. The second stream is slower and source-oriented. Just like OBF, it operates on representations of contents as produced by a given speaker (e.g., *that Tom asserted that p*). In this stream some contents are filtered out because they come from unreliable sources. For example, if Liz knew that Tom is a prolific liar, the content p could be filtered out as being asserted by Tom. Crucially, however, it is not the case that based on the result of the second form of filtering (as OBF and philosophical common view assume) subjects either form a belief that p or no belief at all. The first, content-oriented stream is faster. Therefore, if the content of the assertion passes the gatekeeper of validation, that p is already in the subject’s belief box. Monitoring for the trustworthiness and competence of the source can thus, at most, trigger an attempt of belief revision but does not prevent its formation.

Furthermore, in this model, each stream has its respective real-time filter: validation and source monitoring. Unfortunately, neither of these filters turns out to be effective. Validation filters out only blatantly false and inconsistent information, but it remains virtually helpless against plausible falsehoods (Isberner and Richter, 2014; Marsh et al., 2016). Source oriented filtering, on the other hand, not only has all the problems enumerated by Michaelian and Shieber, but it can at most trigger an attempt to remove already accepted contents from our belief box.

If this model is correct, we go a long way towards demonstrating that: “Humans are constitutionally incapable of learning to respond discriminately to testimony.” (Fricker, 2016, 103) While posing this challenge, Fricker focused on the possibility of the improvement of filtering as it is conceptualized in OBF. But the truth about the mechanisms underlying the comprehension of testimony seems to be more complicated than OBF suggests. There are

good reasons to assume that comprehension consists of two streams and that the faster stream does not involve any source-oriented filtering. Some authors suggest that the form of language processing taking place in the faster stream might be phylogenetically prior, and more directly related to the visual perception (Kissine and Klein, 2013). If this is right, the evolved cognitive setup of our comprehension is not optimized for perfect discrimination between reliable and unreliable testimony, and there is no reason to think that we can alter this predicament if we just try harder.

In this section, I argued that filtering is not effective in the real-time, i.e., the process of filtering, which takes place simultaneously with or right after the process of comprehension, does not allow a particular hearer on a particular occasion to prevent the formation of beliefs based on untrustworthy testimony. Now I proceed to the second part of my discussion on filtering. It is concerned with its long-term effectiveness.

3 Long-term effectiveness of filtering

The discussion about filtering within the epistemology of testimony is predominantly occupied with the matter of its real-time effectiveness. But (i) filtering which we perform is not necessarily limited to the real-time filtering (i.e., filtering taking place upon or right after comprehending given testimony), and (ii) filtering might bring benefits *only* in the long run, i.e., might be long-term effective without being real-time effective.

Both (i) and (ii) are suggested by Dan Sperber (2013) in his response to Michaelian's (2013) critique of Sperber et al.'s (2010) paper about epistemic vigilance.¹⁷ In his (2013), Michaelian deploys a set of arguments similar to these which he offered in (Michaelian, 2010) and which I have summarized

¹⁷Sperber and his colleagues define epistemic vigilance as “a suite of cognitive mechanisms... targeted at the risk of being misinformed by others” (2010, 359). This is exactly what I mean by *the filter* (see footnote 3), thus, I will use *epistemic vigilance* and *the filter* interchangeably.

above, but this time he directs them towards a different opponent. While in (Michaelian, 2010) his target was Fricker’s assumption that filtering is real-time effective and thus provides grounds for a theory of reductive testimonial justification, (Michaelian, 2013) swings at Sperber et al.’s (2010) claim that epistemic vigilance is central for “ensuring reliability and hence the evolutionary stability of communication” (Michaelian, 2013, 37). The evidence from deception detection literature—Michaelian argues—indicates that even if recipients usually avoid being deceived, it is not due to the effectiveness of epistemic vigilance or filtering.

In his response, Sperber (2013) points out that Michaelian’s account of vigilance is too narrow:

Michaelian seems to attribute to us the view that ‘epistemic vigilance is a matter of processes devoted to screening out incoming false information on the basis of available behavioural cues’. Showing that vigilance in this narrow sense is not efficient would, he holds, be quite damaging to our conjecture. This is a misunderstanding. (Sperber, 2013, 65)

This is a misunderstanding, because there is much more to effective filtering than real-time deception detection based on behavioral cues. Not less important are such skills as “drawing inferences from remembered past communicative performance, understanding the sender’s immediate purpose in communicating, or taking into account the sender’s reputation.” (Sperber, 2013, 65).¹⁸ What is, however, central to Sperber’s account, is that epistemic vigilance understood in this broader fashion, is crucial for shaping our social reality. As summarized in the title of his paper: “Speakers are honest because hearers are vigilant.” This hypothesis arises from two independent,

¹⁸Instead of *reputation*, which is a purely social phenomenon, it might be better to talk about psycho-social *demeanor*, i.e., “a constellation of inter-correlated behaviors that function as a gestalt, relating to how people present themselves, the image they convey to others, and how they are perceived by others.” (Levine, 2014, 3).

well-established, and fairly minimal observations. Firstly, *some of us* are *occasionally* successful in detecting dishonesty,¹⁹ either in the real-time or (substantially more often) with a delay.²⁰ Secondly, being caught on dishonesty triggers social retribution (Dunbar, 1996; Dessalles, 2007; Sperber and Baumard, 2012). The combination of these two factors puts pressure on speakers, and results in a high incentive to honesty:

Quasi-universal vigilance makes dishonesty less likely to be beneficial in the short run and more likely to be costly in the long run: falsehoods may be disbelieved, and dishonesty may have reputational costs. (Sperber, 2013, 69)

But maybe speakers are naturally inclined to honesty *independently* of hearer’s vigilance? This seems to be implied by Timothy Levine in his Truth-Default-Theory (TDT) (Levine, 2014, 2019a). According to Levine, it is an empirical fact that dishonesty is not very prevalent. As indicated by survey studies conducted in the US, UK, Netherlands, Japan, and Korea (Serota et al., 2010; Halevy et al., 2014; Serota and Levine, 2015; see also discussion in Levine, 2019a, Chapter 9), “[m]ost communication by most people is honest most of the time” (Levine, 2014, 9) while the majority of lies are produced by a few prolific liars.²¹ Interestingly, Levine explicitly rejects the speculation about an evolutionary arm-race between speaker’s benefiting from deception and hearer’s benefiting from deception detection. Thus, at least seemingly, he contradicts Sperber:

This is where TDT departs from almost all evolutionary thinking about deception. I have heard and read the argument many times

¹⁹Solbu and Frank’s (2019) argue that we might be collectively effective in catching lies thanks to the fact that “some individuals are more apt at being good lie detectors” (p. 40).

²⁰Cf. Park et al. (2002).

²¹Obviously, we have to approach this kind of studies with great caution; they are typically based on self-report surveys conducted within specific, contemporary populations (Sperber, 2013, 67). But there is no reason to cast them aside entirely.

that since humans evolved to deceive, we must have evolved the ability to detect deception. Evolution, it is argued, necessitates a coevolutionary arms race between the ability to deceive and the ability to detect... I do not think that accepting evolution requires accepting a coevolutionary struggle between the ability to deceive and *the ability to detect deception in real time*. (Levine, 2019a, 187, *emphasis mine*)

As we see, Levine rejects the hypothesis that the arms race promotes evolution of the ability to detect deception *in the real-time*. But this is not what Sperber argues for. What he argues for is that the evolutionary arm race forced the receivers to develop a suite of filtering skills, which are effective *in the long-term*: by shaping our social reality into one in which deception is risky and difficult. This much seems to be at least compatible with Levine's own observations:

We have created cultures, religions, and socialization that seek to prevent deception... Prevention is not 100% effective... *Prevention reduces the prevalence and risk of deception to make the truth-default payoff stronger*. It's more efficient to prevent deception than to evolve brains well suited to real-time deception detection. (Levine, 2019a, 189, *emphasis mine*)

In this quote Levine speaks as if prevention was restricted to social conventions and constraints. But if we look under the hood of his theory, we will see that what he postulates is almost indistinguishable from Sperber's vigilance. Based on the evidence about the prevalence of honesty and the ineffectiveness of the real-time deception detection, Levine formulates the core of his Truth-Default-Theory: it is adaptive for participants in communicatory exchanges "to operate on a default presumption that what the other person says is basically honest" (Levine, 2014, 1). A receiver operating under

this presumption remains in what Levine calls a *truth-default state*. However, the truth-default state can be abandoned:

The truth-default state requires a trigger event to abandon it. Trigger events include, but are not limited to (a) a projected motive for deception, (b) behavioral displays associated with dishonest demeanor, (c) a lack of coherence in message content, (d) a lack of correspondence between communication content and some knowledge of reality, or (e) information from a third party warning of potential deception. (Levine, 2014, 9)

Of course, trigger events themselves *have to be detected somehow* to push a receiver out from the default-state.²² Therefore, some kind of “low-key monitoring” (Sperber, 2013, 64) or sensitivity to possible deception has to be active all the time, even in the truth-default state.²³ Despite superficial differences, Levine and Sperber’s theories are compatible. Moreover, they are both simultaneously supported by the evolutionary hypothesis clearly stated by Sperber, as well as the empirical research about the prevalence of honesty collected by Levine.

To sum up. Even though filtering is not real-time effective, it is effective in the long-term. It is not real-time effective because it does not reliably block false contents of comprehended testimonies from falling into our belief box. There are two reasons why this is not the case. First, as argued by Michaelian and Shieber, we are really bad in identifying unreliable testimonies on the

²²A similar picture is suggested by Lipton (2007) in his default-trigger model of testimony. According to Lipton, upon encountering a trigger event, a receiver leaves the default mode and enters the evaluative mode “where he pauses to consider whether he should believe what he has been told” (Lipton, 2007, 241).

²³Obviously, once receivers abandon the truth-default state upon detecting a trigger event, they *do not* become able to effectively detect deception in the real-time. Instead, they become suspicious and motivated to look for further evidence of speaker’s dishonesty. This *might* lead—typically with a delay—to identifying given content as false: “most lies are detected after-the-fact based on either confessions or the discovery of some evidence showing that what was said was false.” (Levine, 2014, 6).

fly. It is, thus, quite uncommon that upon comprehending a testimony we are in possession of reliable information about a speaker's dishonesty, which we could use as a reason not to accept the testimony. Second, as I have argued in Section 2.2 and more extensively in Chapter 2, the very idea that, each time we comprehend a given assertoric utterance, we are free to either accept or reject it, is plausibly based on an idealized and inaccurate picture of linguistic comprehension. If my dual-stream model of comprehension is on the right track, we accept what is said in comprehended assertoric utterances upon only minimal, content-oriented filtering (validation).

Nevertheless, it does not mean that filtering is pointless. Quite the opposite, as I argued in this section, in the long-term filtering makes dishonesty both costly and risky, and thus far less prevalent than it would have been was filtering absent. In the next section, I spell out the consequences this picture has for the debate about the epistemology of testimony.

4 Strong anti-reductionism

Are we by default entitled to believe what we are being told? To answer “no” is to support some sort of reductionism about testimonial entitlement; to answer “yes” is to support some sort of anti-reductionism. In this final section, I will argue for a particular version of anti-reductionism, which results from the above considerations. As I mentioned in the introduction, the very idea that we perform some sort of filtering is not contradictory with anti-reductionism (cf. Goldberg, 2007). Many anti-reductionists suggest that we should distinguish filtering understood as actively looking for reasons to trust a testifier (required by reductionism), from filtering understood as sensitivity to potential signals of dishonesty or incompetence.

A version of anti-reductionism, which ascribes a central role to filtering, is offered by Peter Graham (2010). I have already introduced elements of his theory in Chapter 3, but it is worth reminding some of its crucial aspects here.

Graham formulates his account based on the theory of *etiological functions* (Millikan, 1984). An etiological function of an item is “the effect of the ancestors of the item that explains why the item was replicated” (Graham, 2010, 154). The etiological function of the heart is to pump blood. We have hearts because the effect of pumping blood was beneficial for the systems equipped with organs that were evolutionary ancestors of modern human hearts. Hearts have also other effects, e.g., they make a thumping sound. However, since making the sound *is not* the effect that has let the ancestors of our hearts pass the natural selection, it *is not* the etiological function of the heart of a modern human. Additionally, it is one thing to have an etiological function and another to *function normally*. A heart failure may prevent a particular heart from pumping blood, but the heart does not lose its function of pumping blood. It simply does not function normally, i.e., in accordance with its etiological function.

Graham applies the theory of etiological functions to the problem of epistemic entitlement in general (Graham, 2012) and the problem of entitlement of testimony-based beliefs in particular (Graham, 2010). According to his theory: (i) beliefs formed by normally functioning belief-forming processes that have forming true beliefs reliably as a function enjoy *prima facie pro tanto* entitlement; (ii) comprehension-with-filtering is a belief-forming process that has forming true beliefs reliably as a function. Therefore, testimonial beliefs, which are formed by comprehension-with-filtering of assertoric utterances, enjoy *prima facie pro tanto* entitlement. But why think that (ii) is true? According to Graham, it is an etiological function both of assertion and comprehension of assertoric utterances to induce true hearer beliefs. However, comprehension alone might not induce true hearer’s beliefs *reliably*. This is exactly where filtering helps:

I am happy to grant the possibility that comprehension taken alone, comprehension without filtering, may fail to have inducing true beliefs reliably as a function. . . My thesis is that comprehen-

sion-with-filtering has forming true beliefs reliably as a function, not comprehension neat, comprehension taken alone. *It's the filtering, or so I argue, that is for producing a sufficiently high truth ratio.* (Graham, 2010, 170, *emphasis mine*)

Recently, Mona Simion and Christoph Kelp (2018; cf. Simion, 2020) classified Graham's view as a version of *moderate*, as opposed to *strong*, anti-reductionism. According to strong anti-reductionism, we are *prima facie* entitled to believe whatever we are being told. An *a priori* form of strong reductionism is defended by Tyler Burge (1993; see discussion in Graham, 2018), and an *a posteriori* form of strong anti-reductionism based on the social contract (*Testimonial contractarianism*) is defended by Mona Simion (2020). According to moderate anti-reductionism, on the other hand, some additional condition has to be met for our testimony-based beliefs to enjoy *prima facie* entitlement. In Graham's (2010) theory, this additional condition is that the beliefs have to be produced by the process of comprehension-with-filtering.

What is particularly interesting for my discussion, while criticizing Graham's account, Simion and Kelp appeal to Michaelian's arguments against the real-time effectiveness of filtering (Simion and Kelp, 2018, 2854-5). Graham in fact combines comprehension-with-filtering into a single belief-forming process, and claims that only such a complex process produces true beliefs reliably, but the role played by filtering in Graham's account is in fact much more complex. He never says that filtering is real-time effective in the sense discussed in Section 2 above. He says that "filtering involves sensitivity to counter-considerations: where there are counter considerations of a certain sort, *acceptance would be less likely*" (Graham, 2010, 153, *emphasis mine*). Moreover, Graham is well aware that filtering is long-term effective in the way similar to the one I described in Section 3: "Filtering . . . plays a very important role; it ensures the reliability of beliefs formed on the basis of comprehension, for misleading or untrustworthy assertions often get filtered

out.” (Graham, 2010, 173). If untrustworthy assertions were never filtered out, speakers would lack the incentive to speak truly and, in effect, producing true hearer’s beliefs would not stabilize as the etiological function of the speech act of assertion.

I think that what might be misleading in Graham’s account is the opposition between the process of comprehension taken alone and the process of comprehension-with-filtering. For everyone, (except maybe very small children—I will discuss this matter shortly) there is no comprehension without filtering. Even though they are mostly ineffective in the real-time, some basic filtering mechanisms (such as validation discussed in Section 2.2) always accompany linguistic comprehension. At the same time, we are *prima facie* entitled to believe whatever we are being told, not because we are likely to recognize if an utterance we are about to hear is false, but because we operate in a social environment in which, due to the long-term effectiveness of filtering, dishonesty is risky, costly and thus, not very prevalent. Given that there is no comprehension without filtering, we are *prima facie* entitled to believe whatever we are being told. This is strong, not moderate, anti-reductionism.²⁴

But maybe this clear and harmonious picture is marred if we take into consideration testimonial beliefs acquired by small children? Do comprehension and filtering fall apart in this case? I do not think so. First, if we realize that filtering consists of a multitude of mechanisms: some very sophisticated (like paying attention to someone’s social demeanor), some much more primitive (like detecting incongruences between what one says and our background, e.g., perceptual information), it is quite likely that some basic forms of filtering appear in children’s development simultaneously with the ability to comprehend language. Given that this type of filtering does not have to be real-time effective (to filter does not mean to filter effectively

²⁴It is not obvious to me whether what I present here is an available “strong reading” of Graham’s theory or a distinct theory generated by strengthening Graham’s account.

in the real-time), I do not find this possibility hard to accept. But even if it would turn out that, on some early stages of development, children are already able to comprehend language but lack any sort of filtering mechanisms, their comprehension is a *comprehension-without-filtering* only in the internalist sense. They *themselves* cannot perform the cognitive process of comprehension-with-filtering.²⁵ However, as I have argued throughout this chapter, human communication is never fully bereft of filtering. The social context in which all of us (including small children) operate, is shaped by the long-term effectiveness of filtering. It is sufficient that all adult receivers are, to different degrees, vigilant all the time, for all our beliefs based on the comprehension of testimony to enjoy *prima facie* entitlement.

This view bears important similarities to the way of treating the problem of childhood testimony by Sanford Goldberg (2007).²⁶ Goldberg suggests that, while assessing whether a child acquired knowledge by comprehending a given testimony, we should take into consideration the context in which it happened, especially the role played by the child's adult guardians. According to Goldberg, adult caretakers actively monitor their child's environment and thus "enhance the reliability of a good many of the beliefs that are elicited by the child's encounters with testimony" (2007, 221). While Goldberg's *anti-individualism* relies on *active-monitoring*, the view outlined here can be characterized as a form of *anti-individualism* based on *passive-monitoring*. Taken separately, and from the perspective of a particular instance of reception of testimony, we are all nearly child-like vulnerable to acquiring false beliefs. Our ability to detect deception in the real-time is only slightly better than chance. However, being vigilant by default, every adult language user constantly takes care of the whole linguistic community. Thanks to the long-term effectiveness of filtering, human linguistic communication takes place

²⁵Which might be enough to make the view presented here incompatible with Graham's original account.

²⁶But notice that Goldberg focuses on the acquisition of knowledge, while I focus only on *prima facie* entitlement.

in a sufficiently well-monitored context to make us *prima facie* entitled to believe whatever we are being told.

Finally, I think that the view presented here complements Simion's (2020) *Testimonial contractarianism*. According to *Testimonial contractarianism*, we are *prima facie* entitled to believe whatever we are being told because, in virtue of the social contract in play, speakers are by default compliant with the norms governing speech acts (in particular, the *knowledge norm* of assertion: one should only assert that *p* if one knows that *p*). Speakers are by default compliant with these norms, because this is the rational thing to do if one is not oriented towards one's straightforward and immediate self-interest, but towards long-term, constrained self-interest (Simion, 2020, 23). If we were oriented towards maximizing only straightforward self-interests, it would be rational to lie whenever a given lie might bring about immediate benefits. But since we are oriented towards constrained self-interest, the fact that our lies might be detected (immediately or with a delay; by the receiver themselves or by someone else in the community) changes what is the rational thing to do. Social reputation, which we risk if we are caught on unreliable testimony, is often more valuable than whatever we gain by lying in a particular situation. Thus, what makes it beneficial for speakers to comply with the norms governing speech acts and thus, with a social contract like the one proposed by Simion, is that filtering is long-term effective. The normativity of language use described by Simion seems to be grounded in the psychology and epistemology of the language use described in this chapter.

There is one more, maybe the most radical, conclusion resulting from my discussion in this chapter. Even though currently available empirical data does not allow us to establish it with full certainty, it might be the case that strong anti-reductionism is the only epistemology of testimony that fits our psychological setup. This would be the case if the model of comprehension in which we are always free to choose whether we are going to believe what we are being told (OBF) would turn out to be a false idealization. As I

suggested, it is probable that comprehension of assertoric utterances leads to automatic acceptance of almost everything that is being said; with rejection being available only *ex-post* and requiring extra effort. If this is right, it is good that we are *prima facie* entitled to believe whatever we are being told, because that is exactly what we do.

5 Concluding remarks

So, is filtering effective? In this chapter, I have argued that we should look at this question from a broader perspective. Filtering is not real-time effective because it does not allow us to respond discriminately to particular instances of testimony. We are really bad in online deception detection. But filtering is long-term effective because all (at least all adult) language users are routinely vigilant and being caught on a lie is socially disadvantageous. In result, filtering prevents us from forming a large number of false testimonial beliefs not by turning each of us into a high-functioning polygraph, but by turning the social environment of our communication into one in which such polygraphs are not required.

This way of looking at the effectiveness of filtering allows us to reconsider the role that language comprehension plays in the acquisition of testimony-based beliefs. In particular, it allows us to come to terms with a growing body of empirical research, suggesting that we are not free to either accept or reject whatever we comprehend. It is quite likely that acceptance is the default reaction to comprehended content, and thus real-time effective filtering is simply impossible. Finally, on the ground of the debate about testimonial entitlement, by ascribing the central role to the long-term effectiveness of filtering, these considerations support a form of strong anti-reductionism, i.e., a view according to which we are *prima facie* entitled to believe whatever we are being told.

Chapter 5

Understanding a communicated thought¹

Abstract. The goal of this chapter is twofold. First, we argue that the understanding one has of a proposition or a propositional content of a representational vehicle is a species of what contemporary epistemologists characterise as *objectual understanding*. Second, we demonstrate that even though this type of understanding differs from *linguistic understanding*, in many instances of successful communication, these two types of understanding jointly contribute to *understanding a communicated thought*.

1 Introduction

Grigory Perelman proved the Poincaré conjecture. As a competent English speaker, you understand what you have just read. You have read that Grigory Perelman proved the Poincaré conjecture. We call this kind of understanding,

¹Co-authored with J. Adam Carter and Emma C. Gordon. Material from: Carter, J.A., Gordon, E.C., Grodniewicz, J.P., “Understanding a communicated thought.” *Synthese* (2020), doi.org/10.1007/s11229-020-02854-2 Springer Nature.

however one would like to characterise it in detail,² *linguistic understanding*.

Linguistic understanding of a given utterance differs from understanding of a *proposition* expressed by this utterance.³ To understand the proposition *Grigory Perelman proved the Poincaré conjecture*, it is not sufficient to recognize, that this is what utterances of the English sentence “Grigory Perelman proved the Poincaré conjecture” express in certain contexts. *Prima facie*, it seems plausible that Fields Medal recipient Terence Tao has a rich understanding of the proposition *Grigory Perelman proved the Poincaré conjecture*, while a fourteen-year-old not particularly interested in mathematics—very minimal. Nevertheless, as competent English speakers, they would both agree that this is precisely what is said in the first sentence of the previous paragraph.

There is clearly an *epistemic* difference between understanding a proposition and mere linguistic understanding. What is it? Epistemologists of understanding have by and large been inclined to set this question aside. For example, according to Stephen Grimm (2011), epistemologists seek to uncover the nature of

... understanding of the natural world (broadly understood), and little will be said about how—if at all—the approaches on offer here might relate, for example, to the kind of linguistic understanding we have of concepts or meanings... [because] the way in which we achieve understanding in these areas seems *different enough that it deserves to be dealt with separately* (2011, 84, *our italics*).⁴

²We come back to this question in Section 4.

³We take propositions to be whatever plays the role of contents or objects of attitudes (e.g., belief) and speech acts (e.g., assertion), and semantic values of utterances of sentences in context. We say more on this topic in Section 3.

⁴Grimm uses the term “linguistic understanding” in a very broad sense in which it includes grasping concepts and understanding propositions as meanings of linguistic expressions. We devote Section 4 to argue that the types of understanding we have of propositions and of linguistic entities are distinct.

We disagree. At least, we disagree with the thought that understanding what people tell us (broadly speaking) is interestingly unlike the kind of understanding we have of the natural world, epistemically speaking.

Here is the plan for the chapter. In Section 2, we briefly outline some of the key epistemic features of what epistemologists call *objectual understanding*, e.g., the kind of understanding one might have of a subject matter, such as football or geometry, and how this is typically thought to differ from (mere) propositional knowledge possession.⁵ Next, we show that understanding a proposition is just a special case of *objectual understanding*, viz. the understanding one attains only when one grasps (in an appropriate way) the relations between the constituents of a body of information. In Section 4, we demonstrate how it is that understanding a proposition and linguistic understanding are distinct from one another. Nevertheless, the two types of understanding are closely related; as we argue in Section 5, they jointly, and indispensably, contribute to what we call *understanding a communicated thought*, an important species of understanding that is distinct from, but consists of, both.

2 The epistemology of objectual understanding

A widely held view in epistemology is that *objectual understanding* is a richer cognitive good than mere propositional knowledge.⁶ Objectual understanding attributions take the form “S understands φ ” where φ is (or, can be treated as) as subject matter. “Giles understands algebraic geometry” and “Darla understands football” are paradigmatic sorts of objectual understanding attributions; note that the relations between Giles and algebraic geometry and Darla and football (respectively) are relations between agents and, not

⁵Though see Kelp (2016) for resistance to this orthodoxy.

⁶For instance, Grimm (2011); Elgin (2009); Zagzebski (2001); Greco (2014); Riggs (2003, 2009); Bengson and Moffett (2011) and Kvanvig (2020).

explanations (as in the case of *understanding-why*⁷) but *objects—viz, bodies of information*. The view that objectual understanding can be in some way reduced to knowing an aggregate of propositions is – unlike the corresponding view that understanding-why can be reduced to propositional knowledge – widely rejected,⁸ and there are three key reasons for this: the *argument from luck*, the *argument from degrees*, and the *argument from regress*.⁹

The *argument from luck* goes as follows: Gettier-style cases feature a kind of epistemic luck that is widely taken to be incompatible with propositional knowledge (Pritchard 2005, 2015). If (objectual) understanding is a species of propositional knowledge, we should be able to generate Gettier-style cases for objectual understanding. But we cannot. Therefore, objectual understanding is not a species of propositional knowledge.¹⁰

To use Kvanvig’s (2003) often-cited example here, suppose the subject matter under consideration is the “Comanche dominance of the southern plains of North America from the late seventeenth until the late nineteenth century” (2003, 197). A student checks out and reads a book from the library on this topic, retains the information and can answer relevant questions afterwards, to the extent that it is natural to attribute to her an understanding of the relevant subject matter. Here’s the twist: while each statement in the book read by the student was true, the fact that this is so was the result of a random and fortuitous glitch at the printers, which by dumb luck corrected for all the mistakes. As a result, all the true beliefs formed by reading the book are only accidentally true (e.g., Gettiered). As Kvanvig sees it, this

⁷For a substantial discussion of understanding-why, see Hills (2016).

⁸If explanatory understanding is reducible to propositional knowledge, then Khalifa’s (2013) proposal stands to motivate an avenue for reducing objectual understanding to propositional knowledge by way of reducing objectual understanding to explanatory understanding.

⁹Cf. Kvanvig (2012) for an argument to the effect that understanding, rather than knowledge, has the property of satisfying curiosity.

¹⁰Zagzebski (2001) and Kvanvig (2003) were among the first to suggest that objectual understanding is (unlike propositional knowledge) *compatible* with Gettier-style epistemic luck.

discovery should not lead us to withdraw the initial claim that the student understands the subject matter in question *even though* it should be enough to lead us to deny that the student possesses the relevant items of propositional knowledge. Accordingly, Kvanvig reasons, objectual understanding is (unlike knowledge) compatible with Gettier-style luck¹¹, and thus, the former is not a species of the latter.

Variations on the *argument from degrees* have been widely defended (e.g., Kvanvig, 2003; Pritchard, 2009; Zagzebski, 2001; Riggs, 2003; Grimm, 2012). The core reasoning goes as follows. Objectual understanding is gradient. Whenever one understands something, we can ask to what *degree* they understand it. Propositional knowledge is not gradient in this way. You either know something or you do not: two people can't know that something is true to different degrees. This is so even if the two knowers differ with respect to how well justified they are in believing the target proposition.¹² But if objectual understanding is just a function of knowledge possession—viz., knowledge of a conjunction of propositions—understanding will not be gradient any more than propositional knowledge of a conjunction is gradient.

¹¹As Pritchard (2009) has pointed out, Kvanvig's case becomes even more compelling if we add a twist to the details: suppose Kvanvig's case is modified so that it becomes a barn-façade-style case, where the relevant epistemic luck at play is environmental rather than intervening. To do this, just imagine the book is itself a book each proposition of which the author knew (and so, none of which is Gettierized). But, for the twist, now suppose that this book is surrounded on its shelf by books in the library with inaccurate accounts of the Comanche dominance of the American plains—placed there by a jokester (who overlooked only the genuine book picked by the student). While, à la barn façade cases, propositional knowledge is incompatible with environmental luck—viz., where the belief could easily have been false despite nothing actually going awry (compare: Russell's stopped clock case)—it would seem especially strange to deny that our student understands the Comanche dominance of the American plains simply because he could have easily read a misleading book. Cf. Grimm (2006) for a challenge to the view that understanding differs from knowledge in its resilience to epistemic luck.

¹²See, however, Pavese (2017) for some resistance to this idea in the context of the relationship between knowing how and knowing that. For a more general rejection of this idea, see Hetherington's defence of epistemic gradualism (e.g., Hetherington, 2001, 2011, 2013). For criticism of Hetherington's gradualist position, see e.g., Ginet (2010); Leite (2006); and Faulkner (2003).

Therefore, objectual understanding is not a species of propositional knowledge.¹³

Thirdly, the *argument from regress* proceeds as follows. If possessing objectual understanding is a matter of knowing propositions, then we should expect that, for any (understandable) subject matter (or body of information) φ , there are some n propositions p such that (i) p_1, p_2, \dots, p_n are parts of φ and (ii) S's knowing these n propositions suffices for S's understanding φ . However—and for reasons that parallel the kind of reasoning that features in Lewis Carroll's (1895) “anti-intellectualist” regress¹⁴ (cf., Ryle, 1945)—it remains an open question, for any given number of φ -relevant propositions one knows (and for any further item of propositional knowledge one might add to what one already knows), whether one understands φ . Just as, per Carroll and Ryle, drawing an inference plausibly requires some kind of ability, so does understanding.¹⁵ As Kvanvig (2003, 192) remarks, one can know various items of information but “*understanding is achieved only when [these] informational items are pieced together by the subject in question,*”¹⁶ and the mere possession of such information does not, as this line of argument goes, entail that one is either able or disposed to do this. (Just imagine, for example, a pupil who fails to understand geometry despite coming to know, via reliable testimony with no defeaters, a range of propositions about geometry that a trusted teacher tells him).

What does this ability involve? According to Riggs (2003, 20–21) it involves seeing how the various parts of the understood subject matter “hang together”, something one does only if one “grasps” the relevant relations between the propositions making up the subject matter. Thus, as Grimm

¹³The degree constraint applies primarily to objectual understanding and is less obvious vis-à-vis explanatory understanding (i.e. understanding-why).

¹⁴For some resistance to this line of argument, as well as to Ryle's variation on it in the case of know-how, see Pavese (2015).

¹⁵Elgin (2009) makes a similar point in suggesting that possessing objectual understanding is not simply a matter of believing a long conjunction of relevant propositions.

¹⁶Our italics.

(2014) puts it: to understand one must “be able to see or *grasp* how changes in some of these items will lead (or fail to lead) to changes in the others.”¹⁷ The idea is that S understands φ *only if* S grasps certain relationships between the relevant parts that constitute φ ; Kvanvig and Grimm have referred to what specifically is grasped as *coherence-making relations*, though how to best characterise the grasping requirement remains an issue of some controversy.¹⁸

3 Understanding of a proposition

The matter of how to understand the “grasping” metaphor has been a central issue in work on understanding across disciplines.¹⁹ For one thing, the grasping metaphor features prominently in an issue at the intersection of epistemology, the philosophy of mind and cognitive science—specifically, with regard to what counts as grasping a *concept*. However, the question driving this debate can easily be rephrased as a question about concept *possession*²⁰ as opposed to a question about grasping, (and understanding), *per se*. But what about the debate concerning propositions and propositional contents?

Propositions play a number of central roles in our mental and social lives. They function as:

- Objects or contents of attitudes like belief, speech acts like assertion, and perhaps others in each category;
- (Partial) meanings of utterances of declarative sentences, and perhaps others;

¹⁷See here also Hills (2009) for a similar proposal.

¹⁸For an alternative view, see Kelp (2015), according to whom the understanding-relevant relations between propositions are best characterised as basing relations.

¹⁹For discussion, see Grimm (2011).

²⁰See here, for example, Bealer (1998); Boghossian (2003); Peacocke (1989); Millar (1994).

- Referents of ‘that’-clauses;
- (Primary?) bearers of truth and falsity, and the modalities of truth and falsity: necessity, possibility, probability (subjective or epistemic, objective or metaphysical);
- What gets assessed in determining the validity of arguments.

(García-Carpintero and Jespersen, 2019, 1210)

As we noted at the outset, it is typical of those working on the epistemology of objectual understanding to set the matter of what is involved in understanding a propositions (viz., what play the above roles) aside as though what would be involved—epistemically—in such understanding is fundamentally different from what would be involved epistemically in understanding the “natural world” (Grimm, 2011). But why think this, exactly?

Let us come back to the example we presented briefly in the introduction. Terence Tao and a fourteen-year-old read an utterance of an English sentence “Grigory Perelman proved the Poincaré conjecture.” As a result, they both entertain²¹ the proposition *Grigory Perelman proved the Poincaré conjecture*. Nevertheless, it is at least *prima facie* plausible that Terence Tao’s epistemic relation to this proposition is somehow different than the fourteen-year-old’s (just as Terence Tao’s relation to geometric topology is different than the fourteen-year-old’s). The difference lies in the fact that Tao has a rich degree of understanding of the discussed proposition (and the *whole field of geometric topology*). The fourteen-year-old, on the other hand, has a minimal or no understanding of either of these subject matters.

In this section, we would like to defend two theses:

²¹We use *entertain* here as an umbrella term for whatever type of cognitive attitude one has towards the proposition after comprehending an utterance of a sentence expressing this proposition. Importantly, we *do not* claim that entertaining is a *sui generis* kind of mental attitude.

Weaker thesis: the kind of understanding we aim to have of propositions is similar to the kind of understanding-as-grasping that contemporary epistemologists call objectual understanding.

Stronger thesis: the kind of understanding we aim to have of propositions is type-identical to the kind of understanding-as-grasping that contemporary epistemologists call objectual understanding.

We would like to keep our discussion as neutral as possible with regard to the vast and venerable debate about the nature of propositions. The only assumption we make is that propositions are structured, i.e., that they “are complex entities, entities having parts or constituents, where the constituents are bound together in a certain way” (King 2019).^{22,23}

A convincing argument suggesting that propositions (or at least propositional contents of mental attitudes) are complex and structured entities appeals to their productivity and systematicity.²⁴ Propositional contents are *productive* because anyone with basic conceptual repertoire can, at least in principle, entertain infinitely many of them. If you can entertain a proposition *that the chair stands next to the table*, and you have a concept of SOFA, you can entertain a proposition *that the sofa stands next to the table*, etc.

²²In fact, our account of understanding of a proposition does not require any substantive metaphysical commitment to the “hidden natures” of propositions. It is compatible with deflationary views about propositions according to which either propositions themselves are not structured but they exist only as represented by structured *representational vehicles* (García-Carpintero, *ms*, 3) or “propositions are *abstractions* from (possible) mental state tokens that represent exactly the same” (Grzankowski and Buchanan 2019, 3160) and thus “inherit” the structure from these tokens. If representational vehicles but not propositions are structured, our account would be more precisely characterized as an account of *understanding a propositional content of a vehicle*. Everything we will say can be stated by assuming this more deflationary stance.

²³More concessively still, we commit ourselves neither to a view that propositions are mind-independent entities (such as facts) nor to its negation (see, e.g., Collins, 2018).

²⁴Again, it might be the case that only *representational vehicles* but not propositions are structured and, therefore, that systematicity and productivity should be explained by appeal to the features of vehicles, not propositions (García-Carpintero, *ms*).

Propositional contents are *systematic* because “...our ability to entertain ... one propositional content is intrinsically connected to our ability to entertain other ... propositional contents, so that our ability to entertain the one automatically implies that we can entertain the others” (Duncan, 2018, 353). If you can entertain the proposition *that a chair stands next to the table* you can, by the same token, entertain the proposition *that the table stands next to the chair*.

This said, let’s return to the main question of this section: how should we characterise the epistemic relationship we have to a proposition we take ourselves to understand? We suggest that, given the assumption that propositions are structured, a (singular) *propositional attitude relation* is off the table; after all, if we assume (for *reductio*) that we understand a proposition p just when we stand in some propositional attitude relation to p , we find ourselves left unavoidably with a remainder: namely, what epistemic relation do we have *vis-à-vis* the *constituents* of p ? Recall that productivity and systematicity rationalise that we grasp p by grasping the constituents of p and relations obtaining between them.

In light of this worry, we might be tempted toward an obvious sort of improvement—specifically, we might suppose that we can preserve a propositional-attitude approach *and* (unlike the previous view considered) cover for the remainder that was the cognitive relationship we must bear to the constituents of p . The idea would be something like the following: S understands p just when S (i) bears some propositional attitude relation to p ; *and* (ii) bears some propositional attitude relations $r_1 \dots r_n$ to the *constituents* of p .

An initial reaction here is to consider whether (ii) on this amended view would make (i) redundant. But ultimately, this doesn’t much matter. The problem with the amended view is that there is no suitable way to fill in the details that will not lead the amended view to collapse into the aggregate view we considered in Section 2.

Recall that a fundamental problem for the aggregate view (*vis-à-vis* ob-

jectual understanding in epistemology) highlighted the sense in which objectual understanding appears to have a grasping condition as a necessary condition, and further, that such a condition (as per the regress problem) will not plausibly be satisfied simply by requiring that the agent know an appropriate number of propositions. The subject should also *grasp* the way in which the constituents of the subject matter “hang together.” In the case of propositions it is relatively easy to fill in the details of the relevant grasping condition: grasping the way in which constituents of propositions “hang together” is just grasping the ways in which these proposition are structured, i.e., what are the relations relating their constituents.²⁵

From here it’s not hard to see that understanding a proposition is looking quite a bit like the kind of understanding epistemologists call objectual understanding. Even more, as we will see, understanding of a proposition seems to be compatible with Gettier-style cases that serve to undermine propositional knowledge.

Let’s consider a case structurally similar to Kvanvig’s (2003) library book case (Section 2). Suppose Alex is studying for a biology exam. A disgruntled employee at a company that makes the biology textbooks that Alex’s class is using has tampered with the chapter on the structure and function of the cell, switching around some of the information. Alex, fortunately, grabs a book

²⁵Our story fits particularly well with a solution to the so called *problem of the Unity of propositions* offered by Eklund (2019) (for an exhaustive discussion of the problem see Gaskin (2008)). Here is the problem: “How can there be this complex, the proposition, made up of its constituents, as opposed to merely (the collection of) the constituents themselves?” (Eklund, 2019, 1236). A somehow classical worry is that the problem cannot be solved without falling into a vicious *constitution regress*. Let’s take the proposition that *a* is *F*. We might assume that it has two constituents: *a* and *F*. But what is the difference between this proposition and a mere collection of *a* and *F*? We have to say that in the proposition *a* instantiates *F* and thus that the proposition has a further constituent: the instantiation relation *R*. But in this case, what is the difference between this proposition and a mere collection of: *a*, *F*, and *R*? etc. Eklund’s solution to this problem is to take the relation relating the constituents as *primitive* and not as a further constituent of the proposition: “what accounts for the constituents being combined into the unity they are is how they are related. What would the remaining worry be?” (Eklund, 2019, 1244).

in which the disgruntled employee *attempted* to mix the definitions around but accidentally mixed them back in the original order, leaving the book that Alex grabs with correct descriptions of the elements of a cell. Alex's acquisition of the series of concepts (including NUCLEUS, CYTOPLASM, and MITOCHONDRION) is thus unsafe. Suppose that Alex studies the textbook carefully and is able to use these previously unfamiliar concepts to entertain propositions about cell structure. Moreover, if she hears an utterance of the sentence "Mitochondria contain no DNA" she can tell that the proposition expressed by this sentence is false, just as she can tell that the propositions expressed by the following sentences are true:

- (1) Cytoplasm is enclosed by the cell membrane.
- (2) Eukaryotic cells contain membrane-bound nuclei.
- (3) Mitochondria produce most of the cell energy supply.

In short, and in a familiar sort of way, Alex has come (thanks to her biology textbook) to understand propositions expressed by utterances of sentences (1-3), *even though* the unsafe acquisition of the relevant information undermines her would-be propositional knowledge of the definitions of constituents of the propositions expressed by (1-3).

Taking stock then, understanding of a proposition is, like the kind of objectual understanding that features in mainstream epistemology, (i) not reducible simply to the possession of a propositional attitude or propositional attitudes; (ii) has a grasping condition as a necessary condition; (iii) is compatible with knowledge-undermining epistemic luck. This concludes our argument for the **Weaker thesis**.

Can we go further and establish the **Stronger thesis** as well? If the stronger thesis is right, then understanding a proposition involves a kind of apprehension of the way the propositional constituents stand in relation to each other that is akin to the way understanding a body of information or a subject matter (e.g., algebraic geometry) involves a kind of apprehension of

how the facts constituting this subject matter stand in relation to each other. Here the degree-based *ex-ante* constraint on objectual understanding noted in Section 2 is relevant. If it turns out that, for example, scientific theories, but *not* propositions, can be understood to greater or lesser degrees, this would surely count against the **Stronger thesis**.

As Riggs (2009, 7) notes “*the amount of information* present in someone’s [objectual] understanding can vary.”²⁶ This will be the case, for instance, when we compare the understanding of a housefire possessed by a novice fireman as opposed to the understanding of a housefire possessed by an expert in exothermal reactions. In this respect, objectual understanding can vary along what we might call the *information dimension*. But the degrees of understanding possessed by two individuals might also diverge along what we will call the *action dimension*. As Elgin (2009) notes: “the student who understands geometry can do more with it than the student who just knows all the axioms, the main theorems and their derivations.” **Stronger thesis** (that understanding a proposition is a special case of objectual understanding) predicts, then, that for two individuals, A and B, A’s understanding of some proposition *p* should be able, in principle, to diverge from B’s understanding of *p* along both the informational and action dimensions.

Each of these points can be made rather straightforwardly. Firstly, the point about informational variance: consider that Terence Tao, *as well as* most of his first-year UCLA students, understand the proposition *Grigory Perelman proved the Poincaré conjecture*. Suppose a first-year student at UCLA, Johnny, knows the following things: that the Poincaré conjecture was a problem in geometric topology that is usually illustrated by wrapping a string around a three dimensional object, and that proving the conjecture involved the application of a mathematical method called the Ricci Flow.²⁷ All of this information was gleaned by Johnny (who is only beginning his

²⁶Our italics.

²⁷See Perelman (2002). For an overview of the Ricci Flow, developed by Richard Hamilton, and its applications to the Poincaré conjecture, see Chow, Lu, and Ni (2006).

maths degree) from an article in the *New Yorker*, written for a general audience.²⁸

While Terence Tao and Johnny both understand *Grigory Perelman proved the Poincaré conjecture*, Terence is in a much better position to see how the constituents of this proposition hang together; (after all, Terence, unlike Johnny, has the information necessary to appreciate a more specific sort of thing that *would be involved* in developing a proof of the Poincaré conjecture). Johnny's comparatively impoverished understanding of the proposition is betrayed by the fact that the information he possesses, and which bears on his understanding the proposition, gives him a significantly less rich picture of the subject matter of the proposition *Grigory Perelman proved the Poincaré conjecture*.

Not only can the understanding of a proposition (like understanding a body of information) two individuals possess diverge along an informational dimension (as we saw here) but it can do so as well along what we called the action dimension. Just as, *à la* Elgin (2009), understanding geometry allows one to do more with it than can one who merely knows axioms, main theorems and their derivations, Terence Tao can do more with his understanding of *Grigory Perelman proved the Poincaré conjecture* than Johnny can. Terence's, but not Johnny's, understanding of the proposition that Grigory Perelman proved the Poincaré conjecture includes items of information A, B, and C, where A, B, and C are pieces of information about what Perelman must have proved to have proven the Poincaré conjecture. Terence, but not Johnny, can use A, B and C as premises in his practical reasoning, conditioned on his understanding that Grigory Perelman proved the Poincaré conjecture.

In yet another instance, then, objectual understanding and understanding of a proposition fail to come apart. This concludes our argument for the

²⁸Nasar, S., & Gruber, D. (2006). *Manifold Destiny*, The New Yorker. August, 28, 44-57.

Stronger thesis: the kind of understanding we aim to have of propositions is type-identical to the kind of understanding-as-grasping that contemporary epistemologists call objectual understanding.

4 Linguistic understanding *versus* understanding of a proposition

Since one of the main roles of propositions is “to be the meanings of sentences (in context), or at least the contents expressible by sentences” (Collins, 2018, 3) it is fairly easy to confuse the type of understanding relation we have to propositions with what is commonly discussed in philosophy of language under the label *linguistic understanding*. These two types of understanding are, however, importantly different.

Some popular recent theories identify linguistic understanding of an utterance expressing p with: (i) knowledge that the speaker said that p (Evans, 1982, 311; Heck, 1995, 84), (ii) “conscious awareness” (Hunter, 1998, 560) or “quasi-perception” (Fricker, 2003, 341) that the speaker said that p , (iii) a belief that p (Millikan, 2004, 121; cf. Mandelbaum, 2014), or (iv) a state of entertaining the content p (Longworth, 2018, 822).²⁹

What is common to all these views is that they characterise linguistic understanding as a mental state through which the hearer³⁰ represents the content of a linguistic utterance.³¹ An ability to understand an utterance is thus an ability to recognize what is being said or what is the meaning

²⁹It is not our goal in this chapter to discuss and assess these views. Even if neither of them offers the accurate characterization of the phenomenon of linguistic understanding, they are all on the right track and thus sufficient for our purposes. Limitations of these views are discussed in Chapter 2 of this thesis.

³⁰For convenience, in most our examples we will focus on cases of hearing and linguistic understanding of speech. All that we say is *mutatis mutandis* applicable to linguistic understanding of written word, sign languages such as ASL, etc.

³¹The content is represented either *directly*, i.e., simply as p (Longworth, 2018; Millikan, 2004) or *indirectly*, i.e., as said by a given speaker (Evans 1982; Heck, 1995; Hunter, 1998; Fricker, 2003).

of a given utterance of a sentence in a language one knows. In the case of assertoric utterances, it is an ability to recognize what proposition was expressed through a given utterance.³²

It is, thus, not difficult to demonstrate the difference between linguistic understanding and understanding of a proposition. Let's return again to our example of the proposition *Grigory Perelman proved the Poincaré conjecture* (p). This proposition can be expressed by utterances of sentences of different languages, in particular by an utterance of an English sentence (EN) and a Polish sentence (PL).

(EN) Grigory Perelman proved the Poincaré conjecture.

(PL) Grigorij Perelman udowodnił hipotezę Poincarégo.

To demonstrate that understanding a proposition and linguistic understanding are different phenomena, we will use a comparison class of four protagonists: Terence Tao³³, a monolingual English fourteen-year-old not particularly interested in mathematics, a monolingual Polish mathematician specializing in geometric topology, and a monolingual Polish fourteen-year-old not particularly interested in mathematics.

As we argued in the previous section, both Tao and the Polish topologist understand p . They grasp and can explain what the proof would require, what light it sheds on various related problems in topology, etc. At the same time, Tao lacks linguistic understanding of (PL) while the monolingual Polish mathematician lacks linguistic understanding of (EN). Upon hearing utterances of respective sentences in languages they *do not* speak, they would not realise that what has been just said was p .

³²Crucially, the minimal condition on linguistic understanding is stronger than mere recognition of a linguistic form or recognition of an utterance of a sentence in a given language as an utterance of a sentence in this language. One can, for example, recognize utterances of French sentences (they can tell that someone speaks French) without recognizing what is being said (without linguistic understanding of this utterances).

³³We assume that Terence Tao does not speak Polish. If he does, let us think instead about Tao's counterpart (Tao-minus-Polish) who doesn't.

The situation is different in the case of monolingual teenagers. Let's take as an example the monolingual English teenager, Emily. Emily is capable of linguistic understanding of (EN) but incapable of linguistic understanding of (PL). Utterances of (PL) would sound to her like an incomprehensible babble. Even if she were able to guess that what she heard was an utterance of a sentence in a foreign language (as opposed to, say, gibberish), she would have problems detecting word boundaries, not to mention retrieving meanings of the words uttered or for that matter the whole utterance. Additionally, lacking any competence with geometric topology, Emily has no (or only minimal) understanding of p . Nonetheless, upon hearing an utterance of (EN), she recognises that what has been said was that one Grigory Perelman (whoever that was) proved (whatever that takes) something called "Poincaré conjecture." Obviously, if she did not hear about Poincaré before, she would not recognise his name in the sentence. She might take "poincare" to be the proper name of the conjecture or, if this is the first time she encounters the concept of CONJECTURE, she could have even falsely assume that what Perelman proved was something called "Poincareconjecture." Nevertheless, she could later on use this information to start reading about the mysterious things she has heard about and soon learn a lot about Perelman, Poincaré, and geometric topology. After some time of devoted studies she could acquire a rich understanding of p and the whole field of topology, possibly becoming one of Tao's most promising students.

To sum up. We demonstrated that it is possible to understand a proposition while being unable to linguistically understand utterances of sentences of a given language which express this proposition. It is also possible to linguistically understand utterances of sentences of a given language which express a proposition but fail to understand (or have only a very minimal understanding of) this proposition. This concludes our argument for the divergence between linguistic understanding and understanding of a proposition.

5 Understanding a communicated thought

In Section 3 we have argued that the type of understanding we have of propositions is just a special case of objectual understanding, and in the previous section, that it differs importantly from linguistic understanding. Nevertheless, the fact that these two types of understanding differ does not mean that they do not often co-occur. In fact, cases in which they co-occur are particularly interesting from the point of view of our successful and fruitful communication.³⁴ Arguably, interlocutors communicate most effectively when they not only understand the linguistic expressions used by each other but also understand the things being said, i.e., how the propositions expressed represent the world. We call this *understanding a communicated thought*.³⁵

If u is an utterance of a sentence in a given language, and p is the proposition expressed by u in a given context:

Understanding a communicated thought requires:

- a) linguistic understanding of u , and
- b) (some level of)³⁶ understanding of p

It is easy to apply this template account to the cases of our protagonists from the previous sections. Our u will, again, be either (EN) or (PL), and our p will be *Grigory Perelman proved the Poincaré Conjecture*. Tao will have

³⁴We do not pretend that we provide here sufficient or even necessary conditions for communicative success. All we say is that the type of understanding which we characterize below is among the important factors contributing to communicative success (at least in some situations).

³⁵The phenomenon we characterize below could be more precisely called *understanding of a thought as communicated through a given utterance*. As we have demonstrated above, one could understand a communicated thought (e.g., p) but not *as communicated through a given utterance* (compare Tao's understanding of p as expressed by (EN) but not (PL)). The more accurate name is, however, quite a mouthful, so we will stick to the shorter version: *understanding a communicated thought*.

³⁶As we argued in Section 3, understanding a proposition just like other kinds of objectual understanding, is a gradable matter.

rich understanding of the communicated thought if he hears (EN). He has linguistic understanding of (EN) and rich understanding of the proposition p . However, despite his deep understanding of p , Tao will not understand the communicated thought if he hears (PL), because he does not understand Polish. As we see, the first way in which one may fail to understand the communicated thought is by failing to understand the utterance through which a given thought was communicated.

There is, obviously, another way in which one might fail to understand the communicated thought. It is demonstrated by the example of the English speaking teenager, Emily. Despite her linguistic understanding of (EN), Emily barely understands the thought communicated through the use of (EN) because she barely understands proposition p .

While arguing for the **Stronger thesis** in Section 3, we were comparing Tao's rich understanding of p with his first-year student's (Johnny's) understanding of the same proposition. Our definition accounts for the fact that Johnny's epistemic standing towards p as communicated through an utterance of (EN) is different to both Tao's as well as Emily's. Unlike Emily, Johnny has considerable understanding of the thought communicated through the use of (EN) because he has a considerable degree of understanding of p . Nevertheless, his understanding of a communicated thought is not as rich as Tao's because his understanding of the proposition p is not as rich as Tao's. Since, as we demonstrated in Section 3, the type of understanding we have of propositions is gradable, and since it contributes to the understanding of the communicated thought, the understanding of the communicated thought is itself a gradable matter.

6 Concluding remarks

The goal of this chapter was twofold. First, we argued that the understanding one has of a proposition or a propositional content of a representational

vehicle is a type of what contemporary epistemologists characterize as *objectual understanding*. Second, we demonstrated that even though this type of understanding differs from linguistic understanding, in many instances of successful communication these two types of understanding jointly contribute to *understanding a communicated thought*. At the same time, we think that our discussion makes a case for a closer collaboration between philosophers interested in communication (linguistic communication in particular) and epistemologists. It is through paying attention to both these research fields at the same time that we will be able to tackle the knotty and multidimensional problem of understanding in human interactions.

Summary

In this thesis, I have presented and defended a series of claims regarding the nature and epistemic role of linguistic understanding. Firstly, I have argued that, besides the state- and disposition-sense of “linguistic understanding,” quite commonly discussed in the philosophical debate, there is yet another, often overlooked, process-sense. I have argued that characterizing linguistic understanding as a process is not only justified from the philosophical point of view (linguistic understanding, just like other processes, unfolds over time) but also is very much in line with the current state of the art in empirical language sciences.

Secondly, I have outlined a novel model of the representational structure of linguistic understanding. I have argued that this structure consists of at least three types of interdependent representations generated by a dual-stream process. The model I have offered establishes a middle ground between two popular accounts of the relation between comprehension and acceptance: *Cartesian*, on which we are free to either accept or reject comprehended information, and *Spinozan*, on which we automatically accept everything we comprehend. On my account, we automatically accept everything that passes the content-oriented filter (so-called *validation*), i.e., everything that is not in obvious tension with our easily accessible background knowledge.

Thirdly, I discussed the justification of *comprehension-based beliefs*, i.e., the beliefs about what other people say. I have argued that this justifica-

tion is non-inferential, i.e., that it does not depend on the justification of other beliefs, such as the beliefs about what words the speaker uttered or what sounds they produced. Instead of defending the most common version of non-inferentialism about the justification of comprehension-based beliefs, i.e., a view on which these beliefs are *prima facie* justified by seemings that the speaker said so and so, I have offered a competitive account. On my account, which I call *teleological comprehension-process reliabilism*: (i) beliefs are *prima facie* justified if they are produced by a process that has forming true beliefs reliably as a function, and (ii) language comprehension is a process that has forming true comprehension-based beliefs reliably as a function.

Fourthly, I have argued that despite what is assumed by many participants in the debate, we are not equipped with a mechanism that allows us to react discriminately to particular instances of untrustworthy testimony, i.e., to prevent the formation of beliefs based on such testimony. However, the fact that all, at least all adult members of our linguistic community are vigilant towards the signs of untrustworthiness, and that liars meet social retribution, brings the long-term benefit of decreasing the number of falsehoods and lies we encounter. This account of the psychosocial mechanisms involved in filtering of the comprehended content provides support for the strong anti-reductionism about testimonial entitlement, i.e., the view that we are *prima facie* entitled to believe whatever we are being told.

Finally, together with the coauthors of Chapter 5: J. Adam Carter and Emma C. Gordon, I have argued that *understanding a proposition*, commonly identified with linguistic understanding, is a distinct phenomenon. More specifically, it is a type of objectual understanding, which is gradable, consistent with epistemic luck, and based on a subject's grasping of the coherence-making relation between the elements of a given subject matter. Nevertheless, both linguistic understanding and understanding a proposition play an important role in our everyday communication. In typical cases of

successful linguistic communication, we *understand communicated thought*, i.e., we understand both what proposition has been expressed by the use of a given utterance (linguistic understanding), and this proposition itself.

References

- Adler, J. E. (1994). Testimony, Trust, Knowing:. *Journal of Philosophy*, 91(5):264–275.
- Altmann, G. T. M. and Kamide, Y. (1999). Incremental interpretation at verbs: restricting the domain of subsequent reference. *Cognition*, 73(3):247 – 264.
- Audi, R. (1997). The Place of Testimony in the Fabric of Knowledge and Justification. *American Philosophical Quarterly*, 34(4):405–422. Publisher: University of Illinois Press.
- Audi, R. (2004). The A Priori Authority of Testimony. *Philosophical Issues*, 14(1):18–34.
- Audi, R. (2006). Testimony, Credulity, and Veracity. In Lackey, J. and Sosa, E., editors, *The Epistemology of Testimony*, pages 25–46. Oxford University Press.
- Austin, J. L. (1962). *How to Do Things with Words*. Clarendon Press.
- Balcerak Jackson, B. (2019). Against the perceptual model of utterance comprehension. *Philosophical Studies*, 176(2):387–405.
- Bealer, G. (1998). A theory of concepts and concept possession. *Philosophical Issues*, 9:261–301.
- Beddor, B. (2015). Process reliabilism’s troubles with defeat. *The Philosophical Quarterly*, 65(259):145–159.
- Beddor, B. (*forthcoming*). Reasons for Reliabilism. In Simion, M. and Brown, J., editors, *Reasons, Justification, and Defeat*. Oxford University Press.
- Begg, I. M., Anas, A., and Farinacci, S. (1992). Dissociation of processes in belief: Source recollection, statement familiarity, and the illusion of truth. *Journal of Experimental Psychology: General*, 121(4):446–458.
- Bengson, J. and Moffet, M. (2011). Nonpropositional Intellectualism. In Bengson, J. and Moffet, M., editors, *Knowing How: Essays on Knowledge*,

- Mind, and Action*, pages 161–195. Oxford: Oxford University Press.
- Bever, T. (1970). The Cognitive Basis for Linguistic Structure. In Hayes, J. R., editor, *Cognition and the Development of Language*. Wiley, New York.
- Boeckx, C. and Benítez-Burraco, A. (2014). The shape of the human language-ready brain. *Frontiers in Psychology*, 5:282.
- Boffard, R. (2014, September 22). Can you read a novel in three hours? *the Guardian*. (Active link).
- Boghossian, P. (2003). Blind reasoning. In *Aristotelian Society Supplementary Volume*, volume 77, pages 225–248. Wiley Online Library. Issue: 1.
- Bond, C. F. and DePaulo, B. M. (2006). Accuracy of Deception Judgments. *Personality and Social Psychology Review*, 10(3):214–234.
- Bond, C. F., Kahler, K. N., and Paolicelli, L. M. (1985). The miscommunication of deception: An adaptive perspective. *Journal of Experimental Social Psychology*, 21(4):331–345.
- Bond, Z. S. (2005). Slips of the Ear. In Remez, R. E. and Pisoni, D. B., editors, *The Handbook of Speech Perception*, pages 290–310. John Wiley & Sons, Ltd. Section: 12 _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/9780470757024.ch12>.
- Borsky, S., Tuller, B., and Shapiro, L. P. (1998). “How to milk a coat.” The effects of semantic and acoustic information on phoneme categorization. *The Journal of the Acoustical Society of America*, 103(5):2670–2676.
- Brogaard, B. (2013). Phenomenal Seemings and Sensible Dogmatism. In Tucker, C., editor, *Seemings and Justification*, pages 270–290. Oxford University Press.
- Brogaard, B. (2018). In defense of hearing meanings. *Synthese*, 195(7):2967–2983.
- Brogaard, B. (2019). Seeing and Hearing Meanings: A Non-Inferential Approach to Speech Comprehension. In *Inference and Consciousness*, pages 99–124. Routledge.
- Burge, T. (1993). Content preservation. *Philosophical Review*, 102(4):457–488.
- Burge, T. (1996). Our entitlement to self-knowledge. *Proceedings of the Aristotelian Society*, 96(1):91–116.
- Burge, T. (1997). Interlocution, Perception, and Memory. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 86(1):21–47. Publisher: Springer.

- Campbell, J. (1982). Knowledge and Understanding. *Philosophical Quarterly*, 32(126):17–34.
- Carroll, L. (1895). What the Tortoise said to Achilles. *Mind*, 4(14):278–280.
- Chaiken, S. (1979). Communicator physical attractiveness and persuasion. *Journal of Personality and Social Psychology*, 37(8):1387–1397.
- Chaiken, S. and Eagly, A. H. (1983). Communication modality as a determinant of persuasion: The role of communicator salience. *Journal of Personality and Social Psychology*, 45(2):241–256.
- Chow, B., Lu, P., and Ni, L. (2006). *Hamilton's Ricci flow*. American Mathematical Society.
- Coady, C. A. J. (1973). Testimony and Observation. *American Philosophical Quarterly*, 10(2):149–155. Publisher: [North American Philosophical Publications, University of Illinois Press].
- Coady, C. A. J. (1992). *Testimony: A Philosophical Study*. Oxford University Press.
- Collins, J. (2018). The redundancy of the act. *Synthese*, 195(8):3519–3545.
- Crowther, T. (2009). Watching, Sight, and the Temporal Shape of Perceptual Activity. *Philosophical Review*, 118(1):1–27.
- Crowther, T. (2011). The Matter of Events. *Review of Metaphysics*, 65(1):3–39. Publisher: Philosophy Education Society.
- Crowther, T. (2018). Processes as Continuants and Process as Stuff. In Stout, R., editor, *Process, Action, and Experience*, pages 58–81. Oxford University Press.
- Davies, M. (1989). Tacit Knowledge and Subdoxastic States. In George, A., editor, *Reflections on Chomsky*. Blackwell.
- DePaulo, B. M. and Pfeifer, R. L. (1986). On-the-Job Experience and Skill at Detecting Deception1. *Journal of Applied Social Psychology*, 16(3):249–267.
- Dessalles, J.-L. (2007). *Why we talk: the evolutionary origins of language*. Number 7 in Studies in the evolution of language. Oxford University Press, Oxford. OCLC: 317473148.
- Dijk, T. A. v. and Kintsch, W. (1983). *Strategies of discourse comprehension*. Academic Press, New York.
- Drożdżowicz, A. (2019). Do we hear meanings? – between perception and cognition. *Inquiry*, pages 1–33.
- Dummett, M. A. E. (1993). *The Seas of Language*. Oxford University Press.

- Dunbar, R. I. M. (1996). *Grooming, gossip, and the evolution of language*. Harvard University Press, Cambridge, Mass.
- Duncan, M. (2018). Propositions are not Simple. *Philosophy and Phenomenological Research*, 97(2):351–366.
- Egan, A. (2008). Seeing and believing: perception, belief formation and the divided mind. *Philosophical Studies*, 140(1):47–63.
- Eklund, M. (2019). Regress, unity, facts, and propositions. *Synthese*, 196(4):1225–1247.
- Elgin, C. (2009). Is understanding factive? *Epistemic value*, pages 322–330. Publisher: Oxford University Press Oxford.
- Evans, A. D., Bender, J., and Lee, K. (2016). Can parents detect 8- to 16-year-olds’ lies? Parental biases, confidence, and accuracy. *Journal of Experimental Child Psychology*, 147:152–158.
- Evans, G. (1982). *The Varieties of Reference*. Oxford University Press.
- Falandays, J. B., Batzloff, B. J., Spevack, S. C., and Spivey, M. J. (2018). Interactionism in language: from neural networks to bodies to dyads. *Language, Cognition and Neuroscience*, pages 1–16.
- Faulkner, P. (2003). Good Knowledge, Bad Knowledge. *Mind*, 112(446):346–349. Publisher: Oxford University Press.
- Feeley, T. H. and Young, M. J. (2000). Self-reported cues about deceptive and truthful communication: The effects of cognitive capacity and communicator veracity. *Communication Quarterly*, 48(2):101–119.
- Fodor, J. A. (1983). *The modularity of mind*. MIT press.
- Frazier, L. and Fodor, J. D. (1978). The sausage machine: A new two-stage parsing model. *Cognition*, 6(4):291–325.
- Fricker, E. (1987). The epistemology of testimony. *Proceedings of the Aristotelian Society, supplementary volumes*, 61:57–83. Publisher: JSTOR.
- Fricker, E. (1994). Against Gullibility. In Chakrabarti, A. and Matilal, B. K., editors, *Knowing from Words*. Kluwer Academic Publishers.
- Fricker, E. (1995). Telling and Trusting: Reductionism and Anti-Reductionism in the Epistemology of Testimony. *Mind*, 104(414):393–411.
- Fricker, E. (2003). Understanding and Knowledge of What is Said. In Barber, A., editor, *Epistemology of Language*, pages 325–66. Oxford University Press.
- Fricker, E. (2006). Varieties of Anti-Reductionism About Testimony? A Reply to Goldberg and Henderson. *Philosophy and Phenomenological Research*

- search*, 72(3):618–628.
- Fricker, E. (2016). Unreliable testimony. In McLaughlin, B. P. and Kornblith, H., editors, *Goldman and His Critics*, pages 88–123. John Wiley & Sons, Inc., Hoboken, NJ, USA.
- Friederici, A. D. (2012). The cortical language circuit: from auditory perception to sentence comprehension. *Trends in Cognitive Sciences*, 16(5):262–268.
- Galton, A. (2018). Processes as Patterns of Occurance. In Stout, R., editor, *Process, Action, and Experience*, pages 41–57. Oxford University Press.
- García-Carpintero, M. (*ms*). Pretense, Cancellation, and the Act Theory of Propositions.
- García-Carpintero, M. and Jespersen, B. r. (2019). Introduction: primitivism versus reductionism about the problem of the unity of the proposition. *Synthese*, 196(4):1209–1224.
- Garrod, S. and Pickering, M. J. (2015). The use of content and timing to predict turn transitions. *Frontiers in Psychology*, 6:751.
- Gaskin, R. (2008). *The unity of the proposition*. Oxford University Press, Oxford ; New York.
- Gawronski, B., Sherman, J., and Trope, Y. (2014). Two of What? A Conceptual Analysis of Dual-Process Theories. In Gawronski, B., Sherman, J., and Trope, Y., editors, *Dual-process theories of the social mind*, pages 3–19. Guilford Press, New York.
- Gilbert, D. T. (1991). How mental systems believe. *American Psychologist*, 46(2):107–119.
- Gilbert, D. T., Krull, D. S., and Malone, P. S. (1990). Unbelieving the unbelievable: Some problems in the rejection of false information. *Journal of Personality and Social Psychology*, 59(4):601–613.
- Gilbert, D. T., Tafarodi, R. W., and Malone, P. S. (1993). You can’t not believe everything you read. *Journal of Personality and Social Psychology*, 65(2):221–233.
- Ginet, C. (2010). Réplica a Hetherington. *Veritas (Porto Alegre)*, 55(2).
- Glenberg, A. M. and Kaschak, M. P. (2003). The body’s contribution to language. In *The psychology of learning and motivation: Advances in research and theory, Vol. 43*, pages 93–126. Elsevier Science, New York, NY, US.
- Goldberg, S. C. (2006). Reductionism and the Distinctiveness of Testimonial Knowledge. In Lackey, J. and Sosa, E., editors, *The Epistemology of*

- Testimony*, pages 127–141. Oxford University Press.
- Goldberg, S. C. (2007). *Anti-Individualism: Mind and Language, Knowledge and Justification*. Cambridge University Press.
- Goldberg, S. C. (2010). *Relying on others: an essay in epistemology*. Oxford University Press, Oxford.
- Goldberg, S. C. (2014). Interpersonal Epistemic Entitlements. *Philosophical Issues*, 24(1):159–183.
- Goldberg, S. C. and Henderson, D. (2006). Monitoring and Anti-Reductionism in the Epistemology of Testimony. *Philosophy and Phenomenological Research*, 72(3):600–617.
- Goldman, A. and Beddor, B. (2016). Reliabilist Epistemology. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, winter 2016 edition.
- Goldman, A. I. (1979). What is Justified Belief. In Pappas, G., editor, *Justification and Knowledge*, pages 1–25. D. Reidel.
- Goldman, A. I. (1999). Knowledge in a social world.
- Goldman, A. I. (2008). Immediate Justification and Process Reliabilism. In Smith, Q., editor, *Epistemology New Essays*, pages 63–82. Oxford University Press.
- Graham, P. J. (2006). Liberal Fundamentalism and Its Rivals. In Lackey, J. and Sosa, E., editors, *The Epistemology of Testimony*, pages 93–114. Oxford University Press.
- Graham, P. J. (2010). Testimonial Entitlement and the Function of Comprehension. In Haddock, A., Millar, A., and Pritchard, D., editors, *Social Epistemology*, pages 148–174. Oxford University Press.
- Graham, P. J. (2012). Epistemic Entitlement. *Noûs*, 46(3):449–482.
- Graham, P. J. (2015). Epistemic Normativity and Social Norms. In Henderson, D. K. and Greco, J., editors, *Epistemic Evaluation*, pages 246–273. Oxford University Press.
- Graham, P. J. (2018). Sincerity and the Reliability of Testimony. In Michaelson, E. and Stokke, A., editors, *Lying: Language, Knowledge, Ethics, and Politics*, pages 85–112. Oxford University Press.
- Graham, P. J. and Lyons, J. C. (*forthcoming*). The structure of defeat: Pollock’s evidentialism, Lackey’s distinction, and prospects for reliabilism. In Simion, M. and Brown, J., editors, *Reasons, Justification, and Defeat*. Oxford University Press.

- Greco, J. (2014). Episteme: Knowledge and understanding. *Virtues and their vices*, pages 285–302. Publisher: Oxford University Press Oxford.
- Greco, J. (2015). Testimonial Knowledge and the Flow of Information. In Henderson, D. K. and Greco, J., editors, *Epistemic Evaluation*, pages 274–290. Oxford University Press.
- Grimm, S. (2012). The Value of Understanding. *Philosophy Compass*, 7(2):103–117. Publisher: Wiley-Blackwell.
- Grimm, S. R. (2006). Is understanding a species of knowledge? *The British Journal for the Philosophy of Science*, 57(3):515–535.
- Grimm, S. R. (2011). Understanding. In Pritchard, D. and Bernecker, S., editors, *The Routledge Companion to Epistemology*. Routledge.
- Grimm, S. R. (2014). Understanding as knowledge of causes. In *Virtue epistemology naturalized*, pages 329–345. Springer.
- Grundmann, T. (2009). Reliabilism and the problem of defeaters. *Grazer Philosophische Studien*, 79(1):64–76. Publisher: Brill Rodopi Section: Grazer Philosophische Studien.
- Grzankowski, A. and Buchanan, R. (2019). Propositions on the cheap. *Philosophical Studies*, 176(12):3159–3178.
- Halevy, R., Shalvi, S., and Verschuere, B. (2014). Being Honest About Dishonesty: Correlating Self-Reports and Actual Lying. *Human Communication Research*, 40(1):54–72. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/hcre.12019>.
- Hartwig, M. and Bond, C. F. (2014). Lie Detection from Multiple Cues: A Meta-analysis. *Applied Cognitive Psychology*, 28(5):661–676. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/acp.3052>.
- Hasson, U., Simmons, J. P., and Todorov, A. (2005). Believe It or Not: On the Possibility of Suspending Belief. *Psychological Science*, 16(7):566–571.
- Hauch, V., Sporer, S. L., Michael, S. W., and Meissner, C. A. (2016). Does Training Improve the Detection of Deception? A Meta-Analysis. *Communication Research*, 43(3):283–343.
- Hauser, M., Chomsky, N., and Fitch, T. W. (2002). The Faculty of Language: What is It, Who has It, and How Did It Evolve? *Science*, 298(22):1569–1579.
- Heck, R. (1995). The Sense of Communication. *Mind*, 104(413):79–106.
- Hermes, J., Behne, T., and Rakoczy, H. (2018). The Development of Selective Trust: Prospects for a Dual-Process Account. *Child Development Perspectives*, 12(2):134–138.

- Hetherington, S. C. (2001). *Good knowledge, bad knowledge: on two dogmas of epistemology*. Clarendon Press ; Oxford University Press, Oxford : New York.
- Hetherington, S. C. (2011). *How to know: a practicalist conception of knowledge*. J. Wiley, Chichester, West Sussex, U.K. ; Malden, MA.
- Hetherington, S. C. (2013). Concessive knowledge-attributions: fallibilism and gradualism. *Synthese*, 190(14):2835–2851.
- Higginbotham, J. (1992). Truth and Understanding. *Philosophical Studies*, 65(1-2):3–16.
- Hills, A. (2009). Moral testimony and moral epistemology. *Ethics*, 120(1):94–127.
- Hills, A. (2016). Understanding Why. *Noûs*, 50(4):661–688.
- Huemer, M. (2007). Compassionate Phenomenal Conservatism. *Philosophy and Phenomenological Research*, 74(1):30–55.
- Hume, D. (1975). *David Hume: Enquiries Concerning Human Understanding and Concerning the Principles of Morals (Third Edition)*. Oxford University Press.
- Hunter, D. (1998). Understanding and Belief. *Philosophy and Phenomenological Research*, 58(3):559–580.
- Isberner, M.-B. and Richter, T. (2014). Comprehension and Validation: Separable Stages of Information Processing? A Case for Epistemic Monitoring in Language Comprehension. In Rapp, D. N. and Braash, J. L., editors, *Processing inaccurate information: theoretical and applied perspectives from cognitive science and the educational sciences*, pages 245–276. The MIT Press, Cambridge, Massachusetts.
- Jackendoff, R. (2002). What’s in the Lexicon? In Nootboom, S., Weerman, F., and Wijnen, F., editors, *Storage and Computation in the Language Faculty*, pages 23–58. Springer Netherlands, Dordrecht.
- Jackendoff, R. and Pinker, S. (2005). The Faculty of Language: What’s Special About It? *Cognition*, 95(2):201–236.
- Johnson, M. K., Hashtroudi, S., and Lindsay, D. S. (1993). Source monitoring. *Psychological Bulletin*, 114(1):3–28.
- Johnson-Laird, P. N. (1983). *Mental Models: Towards a Cognitive Science of Language, Inference, and Consciousness*. Harvard University Press, Cambridge, MA, USA.
- Jørgensen, E. (1991). The progressive tenses and the so-called ‘non-conclusive’ verbs. *English Studies*, 72(2):173–182.

- Kaiser, E. and Trueswell, J. C. (2004). The role of discourse context in the processing of a flexible word-order language. *Cognition*, 94(2):113 – 147.
- Kelp, C. (2015). Understanding phenomena. *Synthese*, 192(12):3799–3816.
- Kelp, C. (2016). Towards a Knowledge-Based Account of Understanding. In Grimm, S., Baumberger, C., and Ammon, S., editors, *Explaining Understanding: New Perspectives from Epistemology and Philosophy of Science*, pages 251–269. Taylor & Francis.
- Kendeou, P. (2014). Validation and Comprehension: An Integrated Overview. *Discourse Processes*, 51(1-2):189–200.
- Kendrick, K. H. (2015). The intersection of turn-taking and repair: the timing of other-initiations of repair in conversation. *Frontiers in Psychology*, 6:250.
- Kendrick, K. H. and Torreira, F. (2015). The Timing and Construction of Preference: A Quantitative Study. *Discourse Processes*, 52(4):255–289.
- Khalifa, K. (2013). Is understanding explanatory or objectual? *Synthese*, 190(6):1153–1171.
- King, J. C. (2019). Structured Propositions. In Zalta, E. N., editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, summer 2019 edition.
- Kintsch, W. (1988). The role of knowledge in discourse comprehension: A construction-integration model. *Psychological Review*, 95(2):163–182.
- Kintsch, W. and Mangalath, P. (2011). The Construction of Meaning: Topics in Cognitive Science(2010). *Topics in Cognitive Science*, 3(2):346–370.
- Kintsch, W. and van Dijk, T. A. (1978). Toward a model of text comprehension and production. *Psychological Review*, 85(5):363–394.
- Kissine, M. and Klein, O. (2013). Models of communication, epistemic trust and epistemic vigilance. In Forgas, J. P., Vincze, O., and László, J., editors, *Social cognition and communication*, pages 139–154. New York: Psychology Press.
- Kjelgaard, M. M. and Speer, S. R. (1999). Prosodic facilitation and interference in the resolution of temporary syntactic closure ambiguity. *Journal of Memory and Language*, 40(2):153–194.
- Kraut, R. E. and Poe, D. B. (1980). Behavioral roots of person perception: The deception judgments of customs inspectors and laymen. *Journal of Personality and Social Psychology*, 39(5):784–798.
- Kuperberg, G. R. and Jaeger, T. F. (2016). What do we mean by prediction in language comprehension? *Language, Cognition and Neuroscience*,

- 31(1):32–59.
- Kvanvig, J. (2003). *The Value of Knowledge and the Pursuit of Understanding*. Cambridge University Press.
- Kvanvig, J. L. (2012). Curiosity and a Response-Dependent Account of the Value of Understanding. In Henning, T. and Schweikard, D., editors, *Epistemic Virtues*.
- Kvanvig, J. L. (2020). Understanding. In Aquino, F. and Abraham, D., editors, *Oxford Handbook on the Epistemology of Theology*.
- Köhnken, G. (1987). Training police officers to detect deceptive eyewitness statements: Does it work? *Social Behaviour*. Publisher: John Wiley & Sons.
- Lackey, J. (2006). It Takes Two to Tango: Beyond Reductionism and Non-Reductionism in the Epistemology of Testimony. In Lackey, J. and Sosa, E., editors, *The Epistemology of Testimony*, pages 160–182. Oxford University Press.
- Leite, A. (2006). Epistemic gradualism and ordinary epistemic practice: Response to Hetherington. *Philosophia*, 34(3):311–324. Publisher: Springer.
- Levin, M. (2013). The progressive verb in modern American English. In Aarts, B., Close, J., Leech, G., and Wallis, S., editors, *The Verb Phrase in English*, pages 187–216. Cambridge University Press, Cambridge.
- Levine, T. R. (2010). A few transparent liars. *Communication Yearbook*, 34:40–61.
- Levine, T. R. (2014). Truth-Default Theory (TDT): A Theory of Human Deception and Deception Detection. *Journal of Language and Social Psychology*, 33(4):378–392.
- Levine, T. R. (2019a). *Duped: truth-default theory and the social science of lying and deception*. The University of Alabama Press, Tuscaloosa.
- Levine, T. R. (2019b). An Overview of Detecting Deceptive Communication. In Docan-Morgan, T., editor, *The Palgrave Handbook of Deceptive Communication*, pages 289–301. Springer International Publishing, Cham.
- Levine, T. R. and Knapp, M. L. (2018). Lying and deception in close relationships. In *The Cambridge handbook of personal relationships, 2nd ed.*, pages 329–340. Cambridge University Press, New York, NY, US.
- Levine, T. R., Park, H. S., and McCornack, S. A. (1999). Accuracy in detecting truths and lies: Documenting the “veracity effect”. *Communication Monographs*, 66(2):125–144.
- Levinson, S. C. (2013). Action formation and ascription. In Sidnell, Jack and

- Stivers, T., editors, *The handbook of conversation analysis*, pages 103–130. Wiley-Blackwell.
- Levinson, S. C. (2016). Turn-taking in Human Communication – Origins and Implications for Language Processing. *Trends in Cognitive Sciences*, 20(1):6–14.
- Levinson, S. C. and Torreira, F. (2015). Timing in turn-taking and its implications for processing models of language. *Frontiers in Psychology*, 6:731.
- Lipton, P. (1998). The epistemology of testimony. *Studies in History and Philosophy of Science Part A*, 29(1):1–31.
- Lipton, P. (2007). Alien Abduction: Inference to the Best Explanation and the Management of Testimony. *Episteme*, 4(3):238–251.
- Long, D. L. and Freed, E. M. (2018). A Memory Retrieval View of Text Comprehension. In Millis, K., Long, D. L., Magliano, J. P., and Wiemer, K., editors, *Deep Comprehension*, pages 27–40. Routledge, 1 edition.
- Longworth, G. (2008a). Comprehending speech. *Philosophical Perspectives*, 22(1):339–373.
- Longworth, G. (2008b). Linguistic Understanding and Knowledge. *Nous*, 42(1):50–79.
- Longworth, G. (2010). Some Models of Linguistic Understanding. *The Baltic International Yearbook*, 5(1):7.
- Longworth, G. (2018). Understanding what was said. *Synthese*, 195(2):815–834.
- Lyons, J. (1997). Testimony, induction and folk psychology. *Australasian Journal of Philosophy*, 75(2):163–178.
- Lyons, J. C. (2009). *Perception and basic beliefs: zombies, modules, and the problem of the external world*. Oxford University Press, Oxford ; New York. OCLC: ocn227915785.
- Lyons, J. C. (2015a). Seemings and Justification. *Analysis*, 75(1):153–164.
- Lyons, J. C. (2015b). Unencapsulated Modules and Perceptual Judgment. In Zeimbekis, J. and Raftopoulos, A., editors, *The Cognitive Penetrability of Perception: New Philosophical Perspectives*, pages 103–122. Oxford University Press.
- Lyons, J. C. (2016). Goldman on Evidence and Reliability. In Kornblith, H. and McLaughlin, B., editors, *Goldman and his Critics*. Blackwell.
- Mandelbaum, E. (2014). Thinking is believing. *Inquiry*, 57(1):55–96.
- Mandelbaum, E. and Quilty-Dunn, J. (2015). Believing without Reason:

- or: Why Liberals Shouldn't Watch Fox News. *The Harvard Review of Philosophy*, 22:42–52.
- Marsh, E. J., Cantor, A. D., and M. Brashier, N. (2016). Believing that Humans Swallow Spiders in Their Sleep: False Beliefs as Side Effects of the Processes that Support Accurate Knowledge. In Ross, B. H., editor, *Psychology of Learning and Motivation*, volume 64, pages 93–132. Academic Press.
- Marslen-Wilson, W. and Tyler, L. K. (1981). Central processes in speech understanding. *Philosophical Transactions of the Royal Society of London. B, Biological Sciences*, 295(1077):317–332.
- Mccornack, S. A. and Parks, M. R. (1986). Deception Detection and Relationship Development: The Other Side of Trust. *Annals of the International Communication Association*, 9(1):377–389.
- McNamara, D. S. and Magliano, J. (2009). Toward a Comprehensive Model of Comprehension. In *Psychology of Learning and Motivation*, volume 51, pages 297–384. Elsevier.
- Mercier, H. (2017). How Gullible are We? A Review of the Evidence from Psychology and Social Science. *Review of General Psychology*, 21(2):103–122.
- Michaelian, K. (2010). In defence of gullibility: the epistemology of testimony and the psychology of deception detection. *Synthese*, 176(3):399–427.
- Michaelian, K. (2013). The evolution of testimony: Receiver vigilance, speaker honesty and the reliability of communication. *Episteme*, 10(1):37–59.
- Millar, A. (1994). Possessing Concepts. *Mind*, 103(409):73–73.
- Millikan, R. G. (1984). *Language, Thought and Other Biological Categories*. MIT Press.
- Millikan, R. G. (2004). *Varieties of Meaning: The 2002 Jean Nicod Lectures*. MIT Press.
- Millikan, R. G. (2005). *Language: a biological model*. Clarendon Press ; Oxford University Press, Oxford : Oxford ; New York.
- Mourelatos, A. P. D. (1978). Events, processes, and states. *Linguistics and Philosophy*, 2(3):415–434.
- Nadarevic, L. and Erdfelder, E. (2013). Spinoza's error: memory for truth and falsity. *Memory & Cognition*, 41(2):176–186.
- O'Brien, E. J. and Cook, A. E. (2016). Separating the activation, integration, and validation components of reading. In Ross, B., editor, *Psychology of*

- learning and motivation*, volume 65, pages 249–276. Elsevier.
- O’Shaughnessy, B. (2000). *Consciousness and the world*. Clarendon Press ; Oxford University Press, Oxford : New York.
- Pantazi, M., Kissine, M., and Klein, O. (2018). The Power of the Truth Bias: False Information Affects Memory and Judgment Even in the Absence of Distraction. *Social Cognition*, 36(2):167–198.
- Park, H. S., Levine, T., McCornack, S., Morrison, K., and Ferrara, M. (2002). How people really detect lies. *Communication Monographs*, 69(2):144–157.
- Patterson, K., Nestor, P. J., and Rogers, T. T. (2007). Where do you know what you know? The representation of semantic knowledge in the human brain. *Nature Reviews Neuroscience*, 8(12):976–987.
- Pauker, E., Itzhak, I., Baum, S. R., and Steinhauer, K. (2011). Effects of cooperating and conflicting prosody in spoken English garden path sentences: ERP evidence for the boundary deletion hypothesis. *Journal of Cognitive Neuroscience*, 23(10):2731–2751.
- Pavese, C. (2015). Knowing a rule. *Philosophical Issues*, 25(1):165–188.
- Pavese, C. (2017). Know-how and gradability. *The Philosophical Review*, 126(3):345–383. Publisher: Duke University Press.
- Peacocke, C. (1989). Possession Conditions: A Focal Point for Theories of Concepts. *Mind & Language*, 4(1-2):51–56.
- Perelman, G. (2002). The entropy formula for the Ricci flow and its geometric applications. *arXiv preprint math/0211159*.
- Pettit, D. (2010). On the Epistemology and Psychology of Speech Comprehension. *The Baltic International Yearbook of Cognition, Logic and Communication*, 5(1):9.
- Piest, B. A., Isberner, M.-B., and Richter, T. (2018). Don’t believe everything you hear: Routine validation of audiovisual information in children and adults. *Memory & Cognition*, 46(6):849–863.
- Pinker, S. and Bloom, P. (1990). Natural language and natural selection. *Behavioral and Brain Sciences*, 13(4):707–727.
- Polka, L. and Nazzi, T. (2018). Interacting processes and developmental biases allow learners to crack the “what” code and the “who” code in spoken language. *Applied Psycholinguistics*, 39(4):757–761.
- Pollock, J. (1986). *Contemporary Theories of Knowledge*. Rowman & Littlefield Publishers, Savage, MD.
- Pritchard, D. (2005). *Epistemic Luck*. Oxford University Press.

- Pritchard, D. (2009). Knowledge, Understanding and Epistemic Value. *Royal Institute of Philosophy Supplement*, 64:19–43.
- Pritchard, D. (2015). Anti-luck epistemology and the Gettier problem. *Philosophical Studies*, 172(1):93–111. Publisher: Springer.
- Pryor, J. (2005). There is immediate justification. In Steup, M. and Sosa, E., editors, *Contemporary Debates in Epistemology*, pages 181–202. Blackwell.
- Rayner, K. and Well, A. D. (1996). Effects of contextual constraint on eye movements in reading: A further examination. *Psychonomic Bulletin & Review*, 3(4):504–509.
- Recanati, F. (2002). Does Linguistic Communication Rest on Inference? *Mind and Language*, 17(1&2):105–126.
- Reid, T. (1970). *An inquiry into the human mind*. University of Chicago Press, Chicago.
- Richter, T. (2015). Validation and comprehension of text information: Two sides of the same coin. *Discourse Processes*, 52(5-6):337–355.
- Richter, T., Schroeder, S., and Wöhrmann, B. (2009). You don't have to believe everything you read: Background knowledge permits fast and efficient validation of information. *Journal of Personality and Social Psychology*, 96(3):538–558.
- Richter, T. and Singer, M. (2018). Discourse updating: Acquiring and revising knowledge through discourse. In Schober, M. F., Rapp, D. N., and Britt, M. A., editors, *The Routledge handbook of discourse processes*, pages 167–190. Routledge.
- Riggs, W. D. (2003). Understanding 'Virtue' and the Virtue of Understanding. In DePaul, M. and Zagzebski, L., editors, *Intellectual Virtue: Perspectives From Ethics and Epistemology*, pages 203–226. Oxford: Oxford University Press.
- Riggs, W. D. (2009). Understanding, Knowledge, and the Meno Requirement. In Haddock, A., Millar, A., and Pritchard, D., editors, *Epistemic Value*. Oxford: Oxford University Press.
- Rothstein, S. (2004). *Structuring Events: A Study in the Semantics of Aspect*. Explorations in Semantics. Wiley.
- Rowlands, M. (2010). *The new science of the mind: from extended mind to embodied phenomenology*. MIT Press, Cambridge, Mass. OCLC: ocn502393594.
- Ryle, G. (1945). Knowing How and Knowing That: The Presidential Address. In *Proceedings of the Aristotelian society*, pages 1–16.

- Sacks, H., Schegloff, E. A., and Jefferson, G. (1974). A simplest systematics for the organization of turn-taking for conversation. *Language*, 50(4):696–735.
- Schafer, A. J., Speer, S. R., Warren, P., and White, S. D. (2000). Intonational Disambiguation in Sentence Production and Comprehension. *Journal of Psycholinguistic Research*, 29(2):169–182.
- Schegloff, E. A. (1992). To Searle on conversation: A note in return. In Searle, J. R., Parret, H., and Verschueren, J., editors, *(On) Searle on Conversation*, pages 113–128.
- Schegloff, E. A. (2007). *Sequence Organization in Interaction: A Primer in Conversation Analysis I*. Cambridge: Cambridge University Press.
- Schroeder, S., Richter, T., and Hoever, I. (2008). Getting a picture that is both accurate and stable: Situation models and epistemic validation. *Journal of Memory and Language*, 59(3):237–255.
- Schwanenflugel, P. J. (1986). Completion norms for final words of sentences using a multiple production measure. *Behavior Research Methods*, 18(4):363–371.
- Serota, K., Levine, T., and Boster, F. (2010). The Prevalence of Lying in America: Three Studies of Self-Reported Lies. *Human Communication Research*, 36:2–25.
- Serota, K. B. and Levine, T. R. (2015). A Few Prolific Liars: Variation in the Prevalence of Lying. *Journal of Language and Social Psychology*, 34(2):138–157.
- Shepherd, J. (2018). Intelligent action guidance and the use of mixed representational formats. *Synthese*.
- Shieber, J. (2012). Against Credibility. *Australasian Journal of Philosophy*, 90(1):1–18.
- Shieber, J. (2015). *Testimony: a philosophical introduction*. Routledge, Taylor & Francis Group, New York.
- Simion, M. (2020). Testimonial contractarianism: A knowledge-first social epistemology. *Noûs*.
- Simion, M. and Kelp, C. (2018). How to be an anti-reductionist. *Synthese*, 197(7):2849–2866.
- Singer, M. (2006). Verification of text ideas during reading. *Journal of Memory and Language*, 54(4):574–591.
- Singer, M. (2013). Validation in Reading Comprehension. *Current Directions in Psychological Science*, 22(5):361–366.

- Singer, M. (2019). Challenges in Processes of Validation and Comprehension. *Discourse Processes*, 56(5-6):465–483.
- Sitton, S. C. and Griffin, S. T. (1981). Detection of deception from clients' eye contact patterns. *Journal of Counseling Psychology*, 28(3):269–271.
- Smitterberg, E. (2005). *The progressive in 19th-century English: A process of integration*. Brill | Rodopi.
- Solbu, A. and Frank, M. G. (2019). Lie Catchers: Evolution and Development of Deception in Modern Times. In *The Palgrave Handbook of Deceptive Communication*, pages 41–66. Springer.
- Soteriou, M. (2013). *The mind's construction: the ontology of mind and mental action*. Oxford University Press, Oxford, United Kingdom, first edition edition. OCLC: ocn829743933.
- Soteriou, M. (2018). Experience, Process, Continuity, and Boundary. In Stout, R., editor, *Process, Action, and Experience*, pages 82–101. Oxford University Press.
- Sparks, J. R. and Rapp, D. N. (2011). Readers' reliance on source credibility in the service of comprehension. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 37(1):230–247.
- Sperber, D. (1997). Intuitive and Reflective Beliefs. *Mind & Language*, 12(1):67–83.
- Sperber, D. (2000). Introduction. In Sperber, D., editor, *Metarepresentations: a multidisciplinary perspective*, number v. 10 in Vancouver studies in cognitive science, pages 3–13. Oxford University Press, Oxford ; New York.
- Sperber, D. (2013). Speakers are honest because hearers are vigilant: Reply to Kourken Michaelian. *Episteme*, 10(01):61–71.
- Sperber, D. and Baumard, N. (2012). Moral Reputation: An Evolutionary and Cognitive Perspective. *Mind & Language*, 27(5):495–518. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1111/mila.12000>.
- Sperber, D., Clément, F., Heintz, C., Mascaro, O., Mercier, H., Origg, G., and Wilson, D. (2010). Epistemic Vigilance. *Mind & Language*, 25(4):359–393.
- Steinhauer, K., Alter, K., and Friederici, A. D. (1999). Brain potentials indicate immediate use of prosodic cues in natural speech processing. *Nature neuroscience*, 2(2):191–196.
- Sternglanz, R. W., Morris, W. L., Morrow, M., and Braverman, J. (2019). A Review of Meta-Analyses About Deception Detection. In Docan-Morgan,

- T., editor, *The Palgrave Handbook of Deceptive Communication*, pages 303–326. Springer International Publishing, Cham.
- Steup, M. (2018). Destructive defeat and justificational force: the dialectic of dogmatism, conservatism, and meta-evidentialism. *Synthese*, 195(7):2907–2933.
- Steward, H. (1997). *The Ontology of Mind: Events, Processes, and States*. Oxford University Press.
- Steward, H. (2012). Actions as Processes. *Philosophical Perspectives*, 26(1):373–388.
- Steward, H. (2018). Ocurrent States. In Stout, R., editor, *Process, Action, and Experience*, pages 102–119. Oxford University Press.
- Stivers, T., Enfield, N. J., Brown, P., Englert, C., Hayashi, M., Heinemann, T., Hoymann, G., Rossano, F., de Ruiter, J. P., Yoon, K.-E., and Levinson, S. C. (2009). Universals and cultural variation in turn-taking in conversation. *Proceedings of the National Academy of Sciences*, 106(26):10587–10592.
- Stout, R. (2018). Ballistic Action. In Stout, R., editor, *Process, Action, and Experience*, pages 210–228. Oxford University Press.
- Street, C. N. H. and Richardson, D. C. (2015). Descartes Versus Spinoza: Truth, Uncertainty, and Bias. *Social Cognition*, 33(3):227–239.
- Tanenhaus, M., Spivey-Knowlton, M., Eberhard, K., and Sedivy, J. (1995). Integration of visual and linguistic information in spoken language comprehension. *Science*, 268(5217):1632–1634.
- Vendler, Z. (1957). Verbs and Times. *Philosophical Review*, 66(2):143–160.
- Verkuyl, H. J. (1972). *On the Compositional Nature of the Aspects*. D.Reidel Publishing Company.
- Vrij, A. (2000). *Detecting lies and deceit: the psychology of lying and the implications for professional practice*. Wiley series in the psychology of crime, policing, and law. John Wiley, Chichester ; New York.
- Vrij, A. (2004). Why professionals fail to catch liars and how they can improve. *Legal and Criminological Psychology*, 9(2):159–181.
- Weil, R., Schul, Y., and Mayo, R. (2019). Correction of evident falsehood requires explicit negation. *Journal of Experimental Psychology: General*.
- Wyer, R. S. and Radvansky, G. A. (1999). The comprehension and validation of social information. *Psychological Review*, 106(1):89–118.
- Zagzebski, L. (2001). Recovering Understanding. In Steup, M., editor,

Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue. Oxford University Press.

Zwaan, R. A. (2016). Situation models, mental simulations, and abstract concepts in discourse comprehension. *Psychonomic Bulletin & Review*, 23(4):1028–1034.

Zwaan, R. A., Langston, M. C., and Graesser, A. C. (1995). The Construction of Situation Models in Narrative Comprehension: An Event-Indexing Model. *Psychological Science*, 6(5):292–297.

Zwaan, R. A. and Radvansky, G. A. (1998). Situation models in language comprehension and memory. *Psychological Bulletin*, 123(2):162–185.